

Owner's Manual

Alchemy

PASSPORT[®]



Alchemy

Owner's Manual

Passport Software License

1. Grant of License. In consideration of payment of the License fee, which is a part of the price you paid for this product, PASSPORT, as Licensor, grants to you, the LICENSEE, a nonexclusive right to use and display this copy of Alchemy™ (hereinafter the "SOFTWARE") on a single COMPUTER (i.e., with single CPU) at a single location. PASSPORT reserves all rights not expressly granted to LICENSEE.
2. Ownership of Software. As the LICENSEE, you own the magnetic or other physical media on which the SOFTWARE is originally or subsequently recorded or fixed, but PASSPORT retains title and ownership of the SOFTWARE recorded on the original disk copy and all subsequent copies of the SOFTWARE, regardless of the form or media in or on which the original and other copies may exist. This License is not a sale of the original SOFTWARE or any copy.
3. Copy Restrictions. This SOFTWARE and the accompanying written materials are copyrighted. Unauthorized copying of the SOFTWARE, including SOFTWARE that has been modified, merged, or included with other software, or of the written materials is expressly forbidden. You may be held legally responsible for any copyright infringement that is caused or encouraged by your failure to abide by the terms of this License. This SOFTWARE product is not copy protected, distribution to anyone but yourself is illegal.
4. Use Restrictions. As the LICENSEE, you may physically transfer the SOFTWARE from one computer to another, provided that the SOFTWARE is used on only one computer at a time. You may not electronically transfer the SOFTWARE from one computer to another over a network. You may not distribute copies of the SOFTWARE, or accompanying written materials to others. You may not modify, adapt, translate, reverse engineer, decompile, disassemble or create derivative works based on the SOFTWARE. You may not modify, adapt, translate or create derivative works based on the written materials without the prior written consent of PASSPORT.
5. Transfer Restrictions. This SOFTWARE is licensed only to you, the LICENSEE, and may not be transferred to anyone. In no event may you transfer, assign, rent, lease, sell or otherwise dispose of the SOFTWARE on a temporary or permanent basis except as expressly provided herein.
6. Termination. This License is effective until terminated. This License will terminate automatically without notice from PASSPORT if you fail to comply with any provision of this License. Upon termination you shall destroy the written materials and all copies of the SOFTWARE, including modified copies, if any.
7. Update Policy. Passport may create from time to time, updated versions of the SOFTWARE. At its option, PASSPORT will make such updates available to the LICENSEE who has paid the update fee and returned the Registration Card to PASSPORT. Registered owners of Alchemy are entitled to purchase later revisions of the program for an update fee. Contact: Customer Service, Attention Software Update at the address listed on the back cover for update information.
8. Miscellaneous. This Agreement is governed by the laws of the State of California.

Warranty

THE SOFTWARE AND ACCOMPANYING WRITTEN MATERIALS (INCLUDING INSTRUCTIONS FOR USE) ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND. FURTHER, PASSPORT DOES NOT WARRANT, GUARANTEE OR MAKE ANY REPRESENTATIONS REGARDING THE USE OR THE RESULTS OF THE USE, OF THE SOFTWARE OR WRITTEN MATERIALS IN TERMS OF CORRECTNESS, ACCURACY, RELIABILITY, CURRENTNESS OR OTHERWISE. THE ENTIRE RISK AS TO THE RESULTS AND PERFORMANCE OF THE SOFTWARE IS ASSUMED BY YOU. IF THE SOFTWARE OR WRITTEN MATERIALS ARE DEFECTIVE YOU, AND NOT PASSPORT OR ITS DEALERS, DISTRIBUTORS, AGENTS OR EMPLOYEES, ASSUME THE ENTIRE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

Passport warrants to the original LICENSEE that the disk(s) on which the SOFTWARE is recorded is free from defects in materials and workmanship under normal use and service for a period of ninety (90) days from the date of delivery as evidenced by a copy of the receipt. Further, PASSPORT hereby limits the duration of any implied warranty on the disk to the period stated above. Some states do not allow limitations on duration of an implied warranty, so the above limitation may not apply to you.

PASSPORT's entire liability and your exclusive remedy as to the disk(s) shall be replacement of the disk that does not meet PASSPORT's limited warranty and which is returned to PASSPORT with a copy of the receipt or warranty. If a failure of the disk has resulted from accident, abuse or misapplication, PASSPORT shall have no responsibility to replace the disk(s) or refund the purchase price. Any replacement disk(s) will be warranted for the remainder of the original warranty period or thirty (30) days, whichever is longer.

Defective media returned by the LICENSEE during the warranty period will be replaced without charge, provided the licensee has signed and returned the warranty registration card enclosed with the product, and provided that the returned product has not been subjected to misuse, damage, or excessive wear. Following the initial ninety-day warranty period defective media will be replaced for a replacement fee - contact PASSPORT for details.

All product returns must be accompanied by a Return Authorization Number issued by PASSPORT. Please call PASSPORT and request an RA number before returning any product to PASSPORT. Returned media should be shipped in a protective package (with the RA number on the outside), postage or shipping costs prepaid, to Passport Designs at the address listed below, and accompanied by a brief description of the problem.

THE ABOVE ARE THE ONLY WARRANTIES OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, THAT ARE MADE BY PASSPORT ON THIS PASSPORT PRODUCT. NO ORAL OR WRITTEN INFORMATION OR ADVICE GIVEN BY PASSPORT, ITS DEALERS, DISTRIBUTORS, AGENTS OR EMPLOYEES SHALL CREATE A WARRANTY OR IN ANY WAY INCREASE THE SCOPE OF THIS WARRANTY AND YOU MAY NOT RELY ON ANY SUCH INFORMATION OR ADVICE. THIS WARRANTY GIVES YOU SPECIFIC LEGAL RIGHTS. YOU MAY HAVE OTHER RIGHTS, WHICH VARY FROM STATE TO STATE.

NEITHER PASSPORT NOR ANYONE ELSE WHO HAS BEEN INVOLVED IN THE CREATION, PRODUCTION OR DELIVERY OF THIS PRODUCT SHALL BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL OR INCIDENTAL DAMAGES (INCLUDING DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION, LOSS OF BUSINESS INFORMATION AND THE LIKE) ARISING OUT OF THE USE OF OR INABILITY TO USE SUCH PRODUCT EVEN IF PASSPORT HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. BECAUSE SOME STATES DO NOT ALLOW THE EXCLUSION OR LIMITATION OF LIABILITY FOR CONSEQUENTIAL OR INCIDENTAL DAMAGES, THE ABOVE LIMITATION MAY NOT APPLY TO YOU.

Technical Support

Passport's technical staff are on hand Monday through Friday, 9am to 5pm Pacific time to answer your technical questions. The Tech Support phone number is:

415-726-3826

You can also contact Tech Support via:

- FAX: 415-726-2254, attn: Tech Support
- America On-line: Keyword passport
- CompuServe: GO MIDIBVEN
- Internet: passport@aol.com
- US Mail

For information about new programs and their system requirements, or updates for your existing Passport programs, please call 415-726-0280.

Credits

Alchemy Development Team

Donny Blank, Steven Curtin, Denis Labrecque, Lowell Levinger, Andrew McGowan, Mats Myrberg, Glenn Ruehle, Brent Silveria, Mike Sondin, Richard Vanderlippe, David Willenbrink

Product Manager

John Rezendes

Manual Written by

Josh Rosen

Published by

Passport Designs, Inc.
100 Stone Pine Road
Half Moon Bay, CA 94019

All Manual Contents and Passport Alchemy™ Software Copyright 1995
Passport Designs, Inc.

No part of this manual may be copied, reproduced, translated, or reduced to any electronic medium or machine-readable form without the written consent of Passport Designs, Inc. Passport Designs, Inc. makes no warranties either express or implied with respect to this manual or with respect to the software described in this manual, its quality, performance, merchantability, or fitness for any particular purpose. Passport Designs, Inc. reserves the right to make improvements or changes in the product described in this manual at any time and without notice. All specifications are subject to change without notice.

ALL RIGHTS RESERVED

Table of Contents

Chapter One: About Sound

Introduction	3
Sound Waves and Their	3
Characteristics	3
The Waveform	6
Adding Waves	7
Phase	8
Building a Square Wave	9
Frequency and Harmonics	10
Simple	11
Harmonic	11
Relationships	11
The Harmonic Spectrum	12
Time and Frequency Domains	13
Changing Amplitude Over Time	15
Changing Waveshape Over Time	16
What is Digital Audio?	18
Bit Formats	19
Waveform Sampling	20
Sampling Rate	21
Sampling Rate vs. Sampling Time	23
Aliasing	23
Looping Concepts	25
Short Loops	26
Long Loops	26
Spectrum Analysis: The Fast Fourier Transform	27
FFT	28
Recommended Reading	30

Chapter Two: Guided Tour

Introduction	33
Setting Up Your Network	33
Mini-tour 1: Opening, Editing and Saving a Sound	35
Mini-tour 2: Creating, Looping, and Transferring a Sound	41

Chapter Three: Using Alchemy

Mini-tour 3: Creating a Stereo Sound File	48
Mini-tour 4: Editing a Sound's Harmonic Spectrum	52
Mini-tour 5: Using Alchemy's Enveloping, Time Scaling, and Pitch Shifting	56
Introduction	63
Opening and Saving Files	63
Opening a Standard Sound File	64
Opening a Dyaxis Sound File	65
Opening Multiple Sound Files	65
Playing a Sound File	66
Directly from Disk	66
Saving a New, Untitled Sound File	66
Updating an Existing File with Current Edits	67
Saving a Sound File Under a Different Name or Format	68
DAN: The Distributed Audio Network	68
Adding a Sampler to the Network	69
Deleting a Sampler from the Network	70
Editing a Sampler's Communication Settings	70
Getting a Sound from the Network	71
Getting All Sounds from a Sampler	72
Getting a Waveform Range from the Network	73
Sending Sounds to the Network	74
Replacing an Existing Sound	74
Assigning a New Voice and Key Range	75
Sending All Open Sound Files to a Sampler	77
Sending a Waveform Range	78
Sending Stereo Sounds	78
Visual Editing and Time Domain Processing	79
Waveform Window Display Tools	79
Standard Waveform Window Modes	80
The Selection Mode	80
Selecting a Waveform Range	80
Sliding a Selected Range	81
Extending a Selected Range	82
Creating a Stereo Soundfile	82
Selecting a Single Channel of a Stereo Display	82
Selecting Both Channels of a Stereo Display	83
Selecting a Loop Range	84

Placing an Insertion Point	84
Saving Selection Ranges	84
Recalling Saved Selections	85
Sound File Navigation	86
The Overview Display	86
Creating a Sound File Overview	86
Navigating from the Overview Display	87
Locating in the Overview Display	87
Changing Overview Resolution	87
Locator Icons	88
Centering Your View on Range Start	88
Centering Your View on Range End	89
Centering Your View on Loop Start	89
Centering Your View on Loop End	89
Changing Waveform Resolution (The Zoom Functions)	89
Zooming In on a Waveform	90
Zooming Out on a Waveform	90
Automatically Zooming into a Selected Range	91
Customizing Your Environment	91
Setting Basic Preferences	91
Choosing Axis Units	93
Viewing a Waveform with Axis Markers	93
Setting the Macintosh Sample Playback Rate	93
Editing Memory Use and Splice Options	94
Waveform Editing Functions	95
Cutting a Waveform	95
Copying a Waveform	95
Pasting a Waveform	96
Mixing a Waveform	96
“Out-point” Editing	97
Inserting a Waveform	98
Blending on Edits	98
Clearing a Waveform	99
Extracting a Waveform	100
Undoing a Process	100
Time Domain Processing Functions	101
Fades and Crossfades	101
Setting the Fade Slope	101
Fading In a Waveform	102
Fading out a Waveform	102
Crossfading Two Waveforms	102

Amplitude Scaling	103
Scaling a Waveform	104
Scaling Example	105
Inverting a Waveform	106
Reversing a Waveform	107
Replicating a Waveform	107
Waveform Draw Mode	108
Drawing a Waveform	108
Amplitude and Frequency	109
Enveloping Modes	109
Amplitude Enveloping Mode	110
Frequency Enveloping Mode	111
Editing Envelopes with Pencil and with Knobs	112
Tracing Envelopes	113
A Note on Copying and Pasting Envelopes	113
Amplitude Fit and Amplitude Scale	114
Looping Tools and Methods	114
The Loop Cursors	114
Turning Loops On and Off	115
Changing the Loop Points	115
Loop Splice Mode	116
Using Loop Splice Mode	116
Crossfade Looping	118
Frequency Analysis and Resynthesis	119
The Harmonic Spectrum Display	119
Analyzing a Selected Waveform	120
Selecting Spectra for Editing	120
Selecting an Individual Harmonic	121
Selecting a Group of Harmonics	121
Extending a Selection	122
Editing Harmonic Spectra	122
Editing the Amplitude of a Single Frequency	122
Cutting Frequency Channels	123
Copying Frequency Channels	123
Clearing Frequency Channels	123
Pasting Frequency Channels	124
Mixing Frequency Channels	125
“Out-point” Editing	126
Resynthesizing the Original Waveform	126
Digital Signal Processing Functions	127
Sample Rate Conversion	127
To Resample a Sound	127

Sample Applications

Digital EQ	129
Digitally Equalizing a Waveform	129
Time Scaling and Pitch Shifting	130
Time Scaling a Sound	131
Pitch Shifting a Sound	132
Digital Signal Processing Effects	133
Example: Building a Simple Echo	134
 Application 1	
Looping Workshop	137
Short Loops	137
Long Loops	137
Looping Techniques	138
A Click in the Loop	139
Try it Yourself: Correct a Loop Click	141
Volume Changes in the Loop	141
Tone Changes in the Loop	143
Mirror Splice	144
Splice Loops vs. Crossfade Loops	146
Mirror Crossfade	146
Fade Slopes	147
Loop Start/Loop End Crossfade	149
Other Types of Loops	151
 Application 2	
Building and Using a Universal Sample Library	153
 Application 3	
Stereo Sampling with Psuedo-Stereo Samplers	158
Capturing Stereo Sound	158
Constructing a Stereo Sound File	160
Matching Up the Channels of a Stereo Image	162
 Application 4	
Using Harmonic Analysis/Resynthesis as a Filtering Tool ..	164
 Application 5	
Working with the SP-1200	167
Getting a Sound from the SP-1200	167
Resampling an Existing Sound for the SP-1200	169
Sending a Sound to the SP-1200	170

Chapter Four: Reference

Introduction	177
The Palette	177
Mode Icons	178
Selection Mode	178
Draw Mode	178
Loop Splice Mode	179
Amplitude Enveloping Mode	179
Frequency Enveloping Mode	179
Info Icon (i)	180
Display Icons	180
Speaker Icon	180
Snapshot Icon	180
The Overview Icon	181
The Loop Cursors Icon	181
The Knob/Draw Toggle Icon (enveloping modes active)	181
The Threshold Icon	182
The Axis Markers Icon	182
Process Icons	182
The Fade Out Icon	182
The Crossfade Icon	183
The Trace Envelope Icon (enveloping modes active) ..	183
The Fade In Icon	183
The Invert Icon	183
The Scale Icon	184
The Reverse Icon	184
The Replicate Icon	184
The Amplitude Scale Icon (enveloping modes active) ..	184
The Analyze Icon	185
The Amplitude Fit Icon (enveloping modes active)	185
The Resynthesize Icon	185
The Frequency Mod Icon (enveloping modes active) .	185
The Waveform View Icons	186
The Zoom In Icon	186
The Zoom Out Icon	186
The Fit Selection Icon	186
Cursor Locator Icons	186
Range Start and Range End Select Icons	187
Loop Start and Loop End Select Icons	187
View Memory Buttons	187
Numeric Display Boxes	188

The Windows	188
The Waveform Window	189
Stereo Waveform Display	190
The Enveloping Displays	191
The Harmonic Spectrum Window	193
The Keyboard Dialog	195
The Menus.....	196
The Apple Menu	197
The File Menu	198
New	198
Open...	198
Open Special...	199
Close	199
Save	199
Save As...	200
Revert	200
Import Resource	200
Export Resource	201
Mono to Stereo/Stereo to Mono	201
Soundfile Info...	201
Quit	202
The Edit Menu (waveform window active).....	202
Undo/Redo	203
Cut	203
Copy	203
Paste	203
Mix	204
Insert	205
Extract.....	205
Clear	205
Select All	205
Select Loop	205
Loop Selection	206
Auto Zero	206
Blending.....	206
Edit Options	207
Clear Clipboard	207
The Edit Menu (harmonic spectrum window active).....	207
Undo/Redo	208
Cut	208
Copy	208

Paste	208
Mix	209
Clear	209
Clear Above	209
Clear Below	210
The Edit Menu (enveloping active)	210
Undo/Redo	210
Copy Envelope	210
Paste Envelope	210
Clear Envelope	211
Edit Options	211
Clear Clipboard	211
The Process Menu (waveform active)	211
Fade In	212
Fade Out	212
Xfade	212
Xfade Loop...	213
Invert	213
Scale	213
Reverse	214
Replicate	214
Analyze	214
Resynthesize	214
EQ...	215
Resample...	215
Time Scale...	216
Pitch Shift...	216
The Process Menu (enveloping active)	216
Fade In	217
Fade Out	217
Trace Envelope	217
Invert	217
Scale	217
Reverse	218
Amplitude Scale	218
Amplitude Fit	218
Frequency Mod	219
Frequency Range	219
The Network Menu	220
Get Sound	220
Get Range	221
Get All	221

Send Sound	221
Send Range	222
Send All	222
Instrument	223
New.....	223
Edit... ..	223
Delete	224
Samplers	224
The Windows Menu (waveform active).....	224
Hide/Show Tools	225
Show/Hide Overview	225
Show/Hide Spectrum	225
Tile	225
Strip	226
Stack	226
Waveform	226
Axis Units	226
List of open waveform windows	226
The Windows Menu (harmonic spectrum active).....	227
Harmonic Amp	227
The Action menu	228
Play Sound	228
Take Snapshot	228
Create/Hide Overview	229
Turn Loop On/Off	229
Show/Hide Threshold	229
Show/Hide Rulers	230
Zoom In	230
Zoom Out	230
Full Zoom In	230
Full Zoom Out	231
Fit Selection	231
Audio Output	231
Preferences... ..	231

Appendix

MIDI and SCSI: Communications Standards	235
Ensoniq EPS Specifics	238
E-mu Emax Specifics	239
Akai S900 Specifics	240
E-mu SP1200 Specifics	241

Ensoniq Mirage Specifics	242
CASIO FZ-1/FZ-10 Specifics	243
Roland S-550 Specifics	245
Roland S-50 Specifics	246
Sequential Prophet 2000 and 2002 Specifics	247
Yamaha TX 16W Specifics	248
Oberheim DPX-1 Specifics	249
Ensoniq EPS/SCSI Specifics	250
E-mu EIII Specifics	252
IMS Dyaxis Specifics	255
A Note on MIDI Patch Bays	256

About Sound

Chapter One

Introduction

Alchemy is a sound editing and storage workshop, and as such it can be a very powerful tool for anyone who owns a sampler. By design, Alchemy is simple and logical to use, and it should open some new creative doors. But like any tool, Alchemy is an extension of you; your inspiration is the real creative force. If you have a good basic understanding of how samplers and sound work, you'll be able to push Alchemy to the limits. If you are unfamiliar with the basic concepts behind your technology, you may only be able to scratch the surface of the creative possibilities available to you.

This section of the Alchemy manual has been designed to give newcomers to the sampling world a basic outline of acoustics and sampling technology. Even if you know next to nothing about sampling and sound technology, this section should give you enough information to understand and use your Alchemy program. If you are curious or confused, or even if you just want a quick review, read through this section. The information presented in it may help you get the most out of the program.


Sound Waves and Their Characteristics

Every structure has its foundations. Although the pioneers creating today's most advanced music technologies deal with extremely complex concepts, all of their ideas are constructed of simple building blocks. These blocks are the basics of their science, but they are present in each new development and anyone can understand them. You don't need to know math, physics, or trigonometry. The only real requirements are curiosity and a bit of discipline. The payoff for a short investment of time is a grounded understanding of what's 'going on' inside your technology. This kind of knowledge takes some of the randomness out of creation, and allows you to realize your visions.


The basic building block of all audio technology is the concept of sound. You can think of any sound around you as an event. Something happens, and you perceive it. What's really happening is quite simple, if you remember to think of the air as a substance. To cause a sound you disturb the substance, and the disturbance itself is the sound.

We've all dropped a pebble into a glassy pond and watched as it disturbs the water. Although the pebble drops through the surface, the disturbance it causes sends waves out in every direction. These waves continue outward long after the pebble has disappeared from view. You can picture a sound in exactly the same way. By vibrating something, we create waves in the air which continue outward from the source long after we stop the vibration.


The easiest way to picture this is to look at a simple hypothetical sound machine.



This ideal machine is used to create waves in air just as a wave machine would do it in water. As the motor rotates the arm, the machine repeatedly pushes the air molecules together, creating pressure waves which travel outward from the source. If you could zoom in to look at the waves very closely, you would see areas where the molecules are pushed together and areas where they are much farther from each other.



These squeezed and stretched areas are actually differences in air pressure. The higher pressure areas are above normal (equilibrium) pressure, and the lower pressure areas are below it. These subtle pressure changes are perceived by your ear and translated by your brain into what you would call 'a sound.' One way to represent these pressure changes is with a sine wave.



A sine wave is the simplest representation of a sound, but it's one of the most important. The ideal sound machine you saw earlier would produce an ideal sine wave like this. If you increase the speed of the motor, the machine produces more waves in the same amount of time. We perceive this as an increase in pitch (frequency). If you increase the distance the motor pushes the speaker membrane, the air is pushed farther and the pressure peaks and valleys are made larger. We perceive this as an increase in volume (amplitude). Although this is an oversimplification, it gives you the basic idea.

It may seem strange to concentrate on ideals in such an un-ideal world, but there's a good reason for this. Ultimately, all sounds, no matter how complex, may be broken down into a combination of different sine waves. To understand how this can be true, take a look on the next page to see what happens when sine waves are combined.


The Waveform

When an electronic keyboard instrument is plugged into an amplifier and speaker, it can be used like the ideal machine to create a sound wave. It is communicating electronically with the speaker, rather than mechanically, but can still dictate to the amplifier how the speaker membrane should move. You can use this type of setup to create a relatively accurate version of the ideal sine wave.

Every sine wave has the characteristic 'sine' sound, which is rather smooth and pure. However, simply turning one on and off and varying its pitch might be less than fulfilling. Analog synthesists are familiar with a number of other simple waveforms, including square, triangle, and sawtooth waves. Each of these waveforms has its own characteristic tone or timbre which makes it valuable as a creative tool, and like all sound waves each can be built by adding together a number of sine waves.

Adding Waves


When two waves of any type are added together, the resulting wave can be predicted easily and accurately. The amplitude value of one wave at any point in time is simply added to the amplitude value of the other at that same point in time.



Example A demonstrates how adding two identical waves simply doubles the amplitude at any point. Example B shows that adding a wave to a lower-amplitude inverse of itself evenly decreases the amplitude of the original wave by an amount equal to the amplitude of the inverse wave. Example C illustrates how a wave and its exact inverse cancel each other out. The sine waves in example C are exactly the same in every way, but they are positioned differently in time. This difference in position is called phase.

Phase

The concept of phase is very simple, but it's extremely important to synthesisists. It is part of the basic language of sound science, and you'll run into it often. Phase is a way of describing the positioning of any repeating wave in time. By breaking down one cycle of any periodic (repeating) wave into 360 degrees, like a circle, a frame of reference is created. The degrees are used to describe a certain point or section in the waveform. After going through its full cycle the wave begins to repeat, so the 360th degree is also the 0th degree of the next cycle.




The sine waves in the illustration would sound the same if you listened to them one at a time. However, the effect of adding the same wave to itself in different phases is significant, as we've just seen in the Adding Sine Waves illustration. Adding two identical waves in phase with each other produced the same waveform with twice the amplitude (example A), while adding identical waves 180° out of phase perfectly canceled everything (example C).

Building a Square Wave

Adding sine waves together to produce a complex waveform is called Additive Synthesis. This synthesis method is very popular because it is simple and flexible. Theoretically, all sounds could be constructed using this method, although in practice there are many methods in use.

One way to see how sine waves can do such impressive things is to build a complex wave with them. We'll choose a square wave as our goal and build it in steps.



You can see how the square wave is approximated more closely with every step. Each newly added sine wave flattens the wave's top a bit by pushing up the valleys and pulling down the peaks. At the same time the sides are getting steeper. The added waves are always lined up so that they cross zero when the original wave crosses zero. Since $0 + 0 = 0$, the zero crossings remain unchanged, and the frequency of the wave remains the same. Only the shape evolves.

Unfortunately the square wave is still not finished after adding three waves, hence the 'etc.' Quite a few sine waves need to be added before you would find the square wave to be anywhere near ideal. Still, the experiment illustrates the point that any complex waveform could be built by adding simple sine waves together.


Frequency and Harmonics

In many ways, the square wave experiment offers a unique insight into the nature of sound. It's now very easy to visualize all of the different sine waves, each with a sound of its own, which make up our square wave.

In fact, by looking at the sine waves which make up the square, you can tell something about the timbre the square wave will have. The characteristic sound of each added sine is present in the resultant square wave. To understand this a little better you'll need some basic information about frequency and harmonics.

The concept of frequency is very straightforward. It is described in "cycles per second," which are called *hertz* (abbreviated Hz). The hypothetical sound machine pushes and pulls the speaker membrane through one cycle with every complete revolution of the motor. If the motor spins at 200 revolutions per second, the speaker membrane moves through 200 cycles per second, and a 200 hertz sine tone comes out of the speaker. Since the audible range for humans is generally between 20 hertz and 20,000 hertz, you would hear this as a relatively low sound. When you speed up the motor, more pressure waves are produced per second, and the tone coming out of the speaker is higher.

By looking at a graphic representation of a sine wave, you can figure out its frequency quite easily.




Simple Harmonic Relationships

Since you know that a wave's cycle or period is its full positive and negative swing (one push and pull on the sound machine), you can mark it off on the diagram. By looking at how much time it takes to do this, you can determine the wave's period, and from that you can figure out its frequency. In our example it takes .002 of a second for one cycle. To determine frequency you'll need to figure out how many times the cycle repeats in one second. When you divide 1 second into .002 second chunks, you find that 500 of them will fit. This tells you that the wave will cycle 500 times per second, so the frequency is 500 hertz.

To build a square wave we added a number of sine waves together to approximate what we wanted. If you look back at the Building a Square Wave illustration, you'll see that there is something interesting about the sine waves that were added to the original wave. Three cycles of the first added wave fit in one cycle of the original wave. The second added wave fit exactly five cycles in the same period, and the third added seven. The key word here is exactly.

In the square wave experiment, the added waves were all *harmonics* of the original (fundamental) wave. That means that an exact multiple of each wave fit in the original wave's cycle. A wave which fits two cycles in the original wave's period is called its second harmonic, and the progression continues from there. Here's an illustration.



The Harmonic Spectrum

The frequencies of any wave's harmonics are very easy to calculate, since they are exact multiples of the wave's fundamental frequency. A wave with a frequency of 440 hertz has a second harmonic of 880 hertz, a third harmonic of 1320 hertz, and so on. Theoretically, there is no limit to the number of harmonics a wave can have, but anything over 20,000 hertz is outside of most people's range of hearing.


The square wave you saw earlier was the complex waveform which came from adding together a wave and some of its odd-numbered harmonics. If even-numbered harmonics had been used, a different-sounding complex waveform would have resulted. Therefore the square wave has a particular spectrum which is defined by the harmonics which make it up. This is true for all waveforms.

You've seen how you can use sine waves and their harmonics to create a waveform, but once you've got a waveform, how can you look at it and figure out which sine waves it's made of? Unless you know something ahead of time, looking at the shape of a wave doesn't give very accurate insight into its spectral content.

That's where harmonic analysis comes in. Most musicians have seen a spectrum analyzer at one time or another. Generally these are little boxes with an array of LEDs arranged in vertical lines on the front. Under each vertical line is a frequency number indicating what that LED bank will be showing. Feeding any signal through such an analyzer lights up the LEDs to show how much of each frequency is present. Spectrum analyzers like this are far too inaccurate for our uses and are used for other purposes, but let's assume we have an ideal one.

You play a tape of your sawtooth wave into our ideal spectrum analyzer and it prints out a histogram showing the wave's harmonic content.

Amplitude



This harmonic analysis indicates what sine waves will be needed to make this sound, and how powerful each of them should be. The lowest frequency indicated is the fundamental. All of the rest of the frequencies indicated are harmonics of that fundamental. You can tell this because their frequencies are multiples of the fundamental's frequency. The amplitude at each frequency shows how much of that frequency is present in the analyzed sound. *Note:* The actual amplitude of any harmonic is unimportant. What really matters is the *relationship* between all of the harmonics. Hence we'll use relative "amplitude units" in this example.

You can now sit down at the controls of your ideal synthesizer and build your sawtooth wave. You start with a 100 Hz sine wave that has an amplitude of 1 unit. You add a 200 Hz sine wave with an amplitude of .50 units, then a 300 Hz sine wave with an amplitude of .333..., and so on. With a non-varying tone like a simple sawtooth, a good spectrum analysis is like a recipe. A drawing of a sawtooth waveform alone would have given you much less information about how to build it.


Time and Frequency Domains

So you see how sounds are described in two domains. All visual waveform representations of sound describe it in the *time domain*, because they show the actual amplitude of the waveform at all points in time.

When a sound is analyzed and described in terms of its harmonic content, this is called describing sound in the *frequency domain*, because only the sound's frequencies and their relative powers are shown.

Both of these views have strengths and weaknesses. Time domain representations (waveforms) show what a sound is doing over its full duration, but offer little insight into the dynamic activity of its harmonics during that period. Simple frequency domain representations give much more accurate spectral information, but generally only for an instant in time. They're like snapshots of a moment in the sound's timbral evolution. For sounds such as a simple sawtooth wave (which doesn't really develop over time), a snapshot will often suffice, but for more complex sounds a single spectrum is far too limited.

Perhaps the most flexible way to look at sound is to consider both the time and frequency domains.




Changing Amplitude Over Time

For a simple sawtooth wave, this is all of the information you'll ever need. Unfortunately, our simple sawtooth would be defined as a totally static tone that only exists in two states - on and off, and would probably be more annoying than useful. Complex waveforms in which amplitude and wave shape (spectrum) change over time, are much more dynamic and interesting. Almost all natural sounds, including those produced by acoustic instruments, fall into this category. If you want to understand how these sounds are constructed, a couple of new concepts will be helpful.

When some parameter of a sound changes over time, this change is often described using the concept of *envelope*. One way to visualize a simple envelope is to picture our sawtooth wave coming through a speaker which has a volume control. To make the tone more interesting you might vary its volume over time, alternately turning it up and down.

By changing a tone's volume over time, you are changing its amplitude envelope. This is one extremely simple way to add dynamics to any sound. Most synthesists are familiar with this type of envelope, since it is present in some form on almost all synthesizers. The most common type of amplitude envelope is often called an ADSR, which is short for Attack, Decay, Sustain, and Release. By changing the attack, the time it takes the sound to go from zero volume to full volume is adjusted. Decay defines how quickly and far the volume drops after the attack. The sustain level is the volume at which the sound would remain while you hold down a key. Release is how long it takes the sound's volume to return to zero after you let go of the key.

An ADSR envelope is really just an extreme simplification of how sounds happen in nature. To understand this, look at how the amplitude changes over time in this clarinet waveform:




The inherent limitation of an ADSR envelope is that there are only four points which can be adjusted, and this severely limits your flexibility. The natural waveform above, for example, would require many more than four adjustable points to describe how its volume changes over time.

Changing Waveshape Over Time

The waveform shapes of most natural sounds evolve over time, because the volume (amplitude) of the harmonics that make up those sounds evolve over time. And to complicate matters, the amplitude envelope of every harmonic may be different from that of the other harmonics. As the amplitudes of these harmonics evolve over time, the shape of the wave produced by adding them together also evolves. Look at this test case using an oboe tone.

Amplitude Envelopes of Different Harmonics in an Oboe Tone



As an oboe tone is played, the amplitude envelope of each of its harmonics evolves. As you can see, the amplitudes of the odd harmonics seem to average well above those of the even harmonics. This tells you that the oboe's timbre will have something in common with the square wave's timbre, since both have very strong odd harmonics.

As the oboe's sound evolves over time, the relationships between the harmonics change. Although the odd harmonics were much stronger at the beginning of the tone, their amplitudes decrease so that they are about equal to the even harmonics at the end of the tone. Since the waveform's shape depends on these relationships, it changes as they change.

This complex timbral evolution is one major reason why natural acoustic sounds are so interesting, and so difficult to emulate synthetically. Even to get close to the original oboe sound, a synthesizer would need to have an oscillator and an extremely adjustable envelope generator for each harmonic. The envelope generators in particular present a problem, since they would need hundreds (or thousands) of adjustable points to describe accurately the complex envelopes you saw above.

What is Digital Audio?

This is one reason why sampling technology is so useful. Samplers are able to capture a waveform exactly (or nearly exactly) as it exists. This method manages to avoid the tedium of isolating harmonic content and building sounds one sine wave at a time. The spectrum of the sampled sound and its evolution over time remain intact. Samplers allow you to *begin* with a complex waveform and go from there.

After a sound is sampled, it is possible to analyze and change its time and frequency content in many ways. Although you probably know enough now to begin doing this, a familiarity with sampling technology in general will make it easier.

Digital audio is the process of describing all sound and sound processing by using numbers. Until recently, computers were too slow and memory too costly to make this affordable, but thanks to the technological developments of the last five years, the speed, memory, and cost barriers have been broken. Relatively inexpensive computers are now fast enough to translate sounds immediately into numbers, and to store these numbers for later playback or editing.

Although not everyone prefers digital storage and manipulation of sound to its analog equivalent, the strengths of digitized sound are quite clear. The digital form of sound consists of discrete data and is more flexible, because discrete numbers are much easier to store and maintain than the continuous magnetic charge used in analog audio.

Analog recording, for example, uses a microphone to translate sound into continuous voltage changes. The voltage changes are then fed into a tape deck and recorded as patterns of magnetization in the oxide particles of the recording tape. These patterns are *analogous* to the original waveform. Unfortunately, each time the tape is played back the analogous magnetic information is changed back into voltage information, which degrades it a little. This degradation shows up as ever-increasing audible noise. Also, every time the analog signal is sent somewhere new (when you bounce tracks or add

reverb, for example), the signal is further degraded or distorted by the qualities of the new process or device. Eventually the recording can become so noisy that it can't be used.

With digital recording, the voltage changes coming out of the microphone are discretely measured and translated by an Analog-to-Digital Converter (ADC) directly into numbers. These numbers can then be placed either in the computer's disk memory, or onto another storage medium, such as tape. To play back the recorded sound, the numbers are translated through a Digital-to-Analog Converter (DAC) directly back into the original voltages. These are filtered to remove unwanted high frequencies and then amplified.

Track-bouncing or signal processing can be accomplished by directly changing the numbers which are being used to describe the sound. Since this is done mathematically, little extra noise or distortion is added to the signal. Using the same type of methods, the digital recording can also be edited at will, and with an accuracy which is impossible using analog tape.

So you see that by using a computer to translate sound into digital information, you can avoid degradation of the original, and increase the flexibility of editing. This is essentially what all samplers do, but they do it with different levels of success.

Bit Formats

All forms of sampling machines use some type of internal architecture for storing and manipulating digital data. Perhaps the most fundamental difference between these architectures is the size of a single piece of digital information (often called a word). When you see a machine advertised as 8-bit, 12-bit, 16-bit, or 24-bit, this indicates the size of a single word. The machines are built so that a single word is the standard unit of communication. Bigger words carry more information, and hence allow greater resolution.

You can understand this difference in resolution by thinking of a simple counting system. Between 00 and 10 there are 11 available numbers (including both 00 and 10). Since we are limiting ourselves to two digits, this gives us only 11 possible measurements. We could use this system to indicate the age

Waveform Sampling

of any human by rounding all ages up to the nearest multiple of 10. A 40-year-old would be '4,' a 53-year-old would be '5,' a 76-year-old would be '8.' You can see that this isn't very accurate, but it can still be useful.


However, if your counting system allowed three digits, you would have 101 possible measurements (000 to 100). This makes it possible for you to indicate *exact* age in years, which is obviously much more accurate. Computers with 16- and 24-bit data paths capitalize on this increase in accuracy. With samplers, larger data paths are generally used to increase the accuracy of amplitude measurements, and hence the overall fidelity of the numeric samples.

There are also other ways of increasing sampling accuracy. One of the most popular is called "companding" (from "*compress*" and "*expand*"). Companding samplers use a logarithmic scale instead of a linear one to measure a wave's amplitude. This means that the size of an "amplitude unit" changes according to the range it's in. It sounds complicated, but it really isn't. By using a logarithmic scale to measure amplitude instead of a normal linear scale, lower amplitude values are measured with an increased accuracy, while higher ones lose some accuracy. Since the majority of amplitudes fall in the lower ranges, this yields a general increase in fidelity. Although this method is not without its problems, it does manage to squeeze higher fidelity out of many sampling machines with small word sizes.

To understand more clearly how samplers accomplish the task of creating a digital version of an analog signal, let's look at how a familiar waveform is digitized.


In order to illustrate some basic sampling ideas, we're going to take a look at how a simple sine wave would be sampled. Perhaps the single most important sampling concept is *sampling rate*. You can think of a sample as a single number describing a wave's amplitude at a given time. The sampling rate tells you how many of these samples are taken every second. To see this more clearly, visualize a sine wave.

Sampling Rate




Each of the small square dots along the sine wave is a sample. The amplitude of the original wave has been measured at each of these points to be stored in the computer. You can then use the computer to recall these numbers at a later time and reconstruct the original wave by replaying the samples through a digital-to-analog converter.

The sampling rate you use determines the frequency resolution of your recording. If you take a lot of samples, your data will be much more accurate with respect to the original sound source, but unfortunately more samples require more memory. To see how a lower sampling rate can diminish accuracy, compare the sampled sine wave above to this more complex waveform.



This waveform crosses all of the same points that our sine wave crossed, but you can see that it is much more complicated than the sine. We have made a mistake here by not taking enough samples to represent the wave accurately. The sampling rate is so low that both the sine wave and this much more complex wave would be represented by exactly the same numbers. When this sampled waveform is played back it would sound exactly like the sine wave, because the more erratic, jagged details were not sampled at all.

To remedy this, we could try to sample the complex waveform at a higher rate. In the examples so far, we have been taking 16 samples in every thousandth of a second. This means that we are sampling at a rate of 16,000 samples per second. This is referred to as a sampling rate of 16,000 hertz, or 16 kilohertz (kHz). Lets try doubling this to a rate of 32 kHz.



You can see that the black dots representing our samples describe the complex waveform much more accurately when the sampling rate is increased. When this sampled waveform is played back, its shape (and sound) would differ quite noticeably from that of the sine wave. If you increased the sampling rate again, the resulting sampled waveform would be even more similar to the original.


Sampling Rate vs. Sampling Time

Why not sample everything at the highest possible rate? There's a good reason. Any sampling machine has an upper limit to the number of total samples it can store. This limit is usually expressed as a certain time (in seconds) at a given sampling rate. For example, a sampling keyboard might have 10 seconds of sampling time at a sampling rate of 50 kHz. This tells you that it can store a total of 500,000 samples (10 seconds multiplied by 50,000 samples per second). What if you wanted to sample a 20 second orchestral finale?

You know that you have 500,000 samples to work with. If you want to increase your overall sampling time, you can do it by decreasing your sampling rate. If you only take 25,000 samples per second, for example, it would require 20 seconds to use all of your 500,000 total samples. However, this isn't without its cost. As you saw above, lower sampling rates have a lower resolution and are therefore less accurate. This decrease in accuracy often shows up as a 'graininess,' much like a low resolution photograph. Lower sampling rates also cause other, much more serious problems.

Aliasing

Without question, one of the major problems you will encounter when sampling is aliasing, which is also called low-frequency foldover. The basic rule of aliasing is that any frequency which is above half of the sampling rate will generate a potentially unpleasant and unwanted lower frequency "alias" of itself. You can figure out the foldover frequency in most circumstances by subtracting the output frequency from the sample rate. For example, if your sampling rate is 20 kHz, playing back a sampled frequency of 15 kHz will fold over a new and possibly irritating 5 kHz frequency. The reason for aliasing is quite simple.



As you can see in the diagram, the samples (indicated by black dots) of the original wave actually describe at least two different waves. Since the sampling rate was too low in relation to the frequency of the original wave, trying to recreate that wave causes an ambiguity which results in the generation of both the original wave, and a new, lower tone.



Note: An interesting example of aliasing is the “wagon wheel” effect you’ve probably noticed in film. Even though a wagon is moving and you can see the wheels turning, the spokes appear to be stationary (or going backwards). If you assume that 25 film frames are being taken every second, and that the wagon wheel is turning at 25 revolutions per second, you’ll see the problem. Every film frame catches the wheel in the same position, although it also captures some of the blur of its motion. The non-moving wheel image is an alias of the actual moving wheel. Each film frame is a “visual sample,” and our sampling rate is simply not high enough to avoid the creation of a confusing image which is not present in reality.

In sampling the wave as we did, we created our own “wagon wheel” effect. The wave we were trying to sample repeats just over three times per thousandth of a second, which gives it a frequency above 3000 Hz. We have been sampling this sine

Looping Concepts

wave exactly six times per thousandths of a second, which means our sampling rate is *exactly* 6000 Hz. Although the frequency we are sampling is only slightly over half of the sampling rate, the alias it causes is very noticeable. The only remedy for this problem is to make sure that your sampling rate is at least twice the highest frequency you are sampling. If you can't adjust your sampling rate, then you have to adjust the sound you are sampling. This can be done by filtering out all frequencies above half of the sampling rate, which is usually accomplished using an *input sampling filter*.

The highest frequency which can be sampled at a given sampling rate is called its Nyquist frequency. It is always equal to half of the sampling rate. If you want to create clean samples which are harmonically true to their source, you must observe this limit.

By nature, there is a direct relationship between sampling capabilities and computer memory. The more available memory you have, the more sampled sounds you can keep at your finger tips.

But more memory costs more money. For this reason, musicians generally want to squeeze the most out of the memory they have, and this is the origin of looping.

The technique of looping is based on the fact that pitched sound is repetitive in nature. If you think of our original sine wave, the first cycle looked exactly like the second, third, fourth, and so on. If you wanted to play ten seconds of a sine wave from a sampling keyboard, you could sample ten seconds of a sine wave from an analog synthesizer, and simply play it back. If your keyboard was capable of 10 total seconds of sampling at your selected rate, it would now be full, and you would have no room for other sounds.

Short Loops

You could avoid this waste of memory by creating a short loop. Instead of sampling ten seconds of a sine wave, you could simply sample one cycle. You lose nothing by doing this because each of the sine wave's cycles is the same. You can then instruct your sampler to play the single cycle over and over. Each time the end of the waveform is reached, your sampler jumps back to the beginning and starts over. This "looping" continues indefinitely (for as long as you hold down a key).

Compared to our first method, this short loop will only require a tiny fraction of a second of sampling time, which leaves you well over nine seconds of valuable time to sample other sounds. And best of all, you can play it for as long as you like.

Long Loops

Although a sine wave is an extremely simple waveform, short loops will work also well with any truly repeating waveforms, such as square, triangle, or sawtooth waves. However, the whole business of looping becomes much more challenging when you're dealing with more complex waveforms.

Imagine the waveform of an entire orchestra, which changes shape dramatically over time. Repeating a single wave cycle would accomplish almost nothing in terms of maintaining the harmonic evolution of the original tone. In such cases, long loops can be very useful.

Long loops take advantage of the fact that most natural sounds have very dynamic beginnings, but stabilize considerably over time. Although this isn't really true for maracas or a whip-crack, it often is for a cello, orchestra, or human voice. Let's assume you sample your own voice singing 'tah...' for approximately one second. Every time you play this sample back, it will last approximately one second—although it will be shorter at higher pitches and longer at lower ones. If you would like the pitch to continue for as long as you hold down the key, you'll need to build a long loop.

Spectrum Analysis: The Fast Fourier Transform

You can think of the long loop of your voice like this: When you press a key, the keyboard will play back the full 'tah...', and when it reaches the end it dives back in at the 'ah..' part of the sound. This 'ah...' loop would then repeat until you stop pressing the key. The whole process gives the impression of unlimited sampling time, while really only using up one second.

This analogy makes long loops seem a little simpler than they are. Much care must be taken in matching the volume and spectrum of the loop end with that of the loop start. Remember that the sound reaches the loop's end and is then replayed from an earlier point. When you play back a looped sound, your sampler jumps immediately from the loop's end to the loop's start, so the wave shape and position at the loop's end must closely match the wave shape and position at the loop's start.

This is not always easy, and there are a number of tricks for accomplishing it. To learn more about understanding and constructing loops, see the Looping Workshop in the Applications at the end of the Using Alchemy section.

Spectral content lies at the base of the most difficult looping problems. The problem is this: A waveform is a time domain representation of a sound, so you can't be sure of its exact spectral content. From a waveform, you can see how long a sound is, how loud it is, and that it changes timbrally over time. An accurate view of *how* it changes timbrally is not evident. In loops, we avoided this problem by using mirror splices and crossfades. We couldn't change the spectral content of our splice points, so we cut in (and mixed in) waveforms with the spectral qualities we wanted. While you were doing this, you might have asked, "Why not just adjust the spectral content of the waveform that's there, instead of replacing it with something else?"

FFT

There is a reason. There's simply no accurate way to adjust complex spectral information in the time domain. To change the spectrum of a waveform in a precise fashion, you might try to hand draw new waveforms over and over until you got closer to your desired tone. But even if you knew the harmonics which you wanted in the final waveform, you would have little information about how a waveform containing those harmonics would be shaped.

The Fast Fourier Transform (FFT) offers a way out of this dead end. If you remember the earlier discussion of the time and frequency domains, you'll recall the ideal spectrum analyzer. We fed a waveform into it, and it gave us a display of that waveform's spectrum. What if we could make this machine go backwards? What if we could feed a list detailing the spectral content of a sound into it and have it produce that sound (and a drawing of the its waveform)? The FFT allows us to do this.

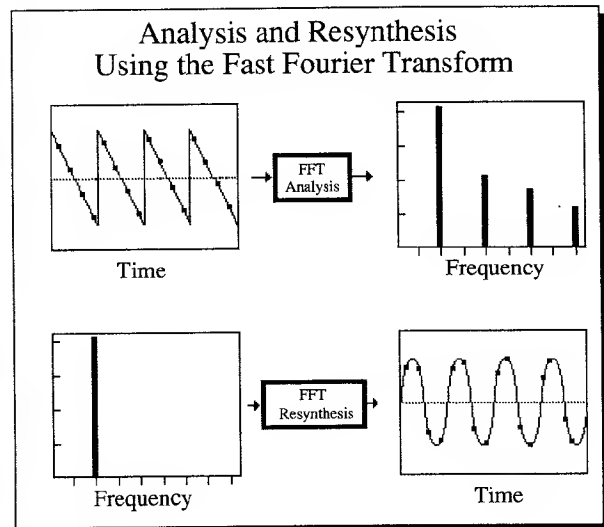
The FFT is a mathematical procedure which can be used to analyze any sampled waveform in order to determine the frequencies, relative amplitudes, and phase relationships of its harmonics. The levels of these harmonics can then be manipulated, and the FFT can be reversed to resynthesize the new waveform. You can think of the FFT as a translator between the time and frequency worlds.

You can use the FFT to produce a detailed spectrum analysis of almost any sized waveform, but it may not work exactly the way you expect. For example, if you analyze a waveform which consists of 1024 samples, the FFT breaks this down into 512 different frequencies. The number of frequencies in the final analysis is always equal to half of the number of samples analyzed. When a larger waveform (consisting of more samples) is analyzed, its harmonic content is always broken down into more frequencies than a smaller waveform.

This may sound confusing, but it makes sense. The reason for it is found in the Transform's mathematics, which are relatively complex. The Transform does its work by analyzing all

of the selected samples in a waveform and using integral calculus on some simple trigonometric relationships to figure out what sine waves make up the waveform, and what their amplitudes are. These sine waves are the harmonics. When long waveforms consisting of many samples are analyzed, there are more possible sine waves which could be present. After analysis is over, one frequency band is displayed for each possible sine wave. With the Fast Fourier Transform, a sine wave can be defined by as few as two samples. Hence there are always half as many possible sine waves as there are samples analyzed. Luckily, most sounds which we call “musical” are not entirely random, and consist of predictable harmonic series. Because of this, the majority of displayed frequency bands for large waveforms will usually be empty.

To illustrate how the FFT works, let's look at an example. First, we'll analyze the ideal sawtooth we looked at earlier in the section. Then we'll change its spectral content, and resynthesize the new waveform.



Although the sampling rate is too low for a true application, this FFT analysis/resynthesis experiment illustrates what the

FFT is all about. The black dots show that we will be analyzing 16 samples which make up a sawtooth wave, so we know that we will have eight frequency bands in the spectrum display. Notice that four of the bands are empty. Now we cut out all of the indicated frequencies except the fundamental (the largest), and resynthesize.

You shouldn't be surprised by what you have. You'll remember that each harmonic may be represented as a sine wave, so resynthesizing any single harmonic will produce a sine wave which corresponds to its frequency. By completing this analysis/resynthesis, we have accomplished what we started out to do. We have changed the spectral content of a sampled waveform with great precision, and then we used the inverse FFT to resynthesize the new wave's shape.

Now you have the basic information you need to understand Alchemy. The Guided Tour in the next chapter will give you some ideas about what Alchemy can do.

Recommended Reading

Introduction to Computer Music, by Wayne A. Bateman, John Wiley & Sons, 1980

Principles of Digital Audio, by Ken C. Pohlman, Howard W. Sams & Co., 1985

Musical Applications of Microprocessors, by Hal Chamberlin, Second Edition, Hayden Books, 1985, 87

The Technology of Computer Music, by M. V. Mathews, The M.I.T. Press, 1969

Chapter Two

Guided Tour

Introduction

This is your Alchemy guided tour, and it will be your first hands-on experience with the program. It has been designed to step you through some of the most important Alchemy functions in order to illustrate what the program does, and how it does it. To go through this tour you'll need your Macintosh, a MIDI interface, a sampler, and the program. You don't need to know anything to complete the tour, just follow the directions, watch, and learn. Each step in any process will be listed in bold print and then explained.

While you are touring Alchemy, you'll come up with some questions about how to do things not covered in the directions. It may help you to learn the program if you keep a list of your questions and pursue them further after you complete the tour. The guided tour is a quick means of introducing you to Alchemy, and is not meant to teach you the program or provide reference information.

Most of the mini-tours in the Guided Tour make extensive

use of the view memories. Each time you click on a view memory number, you recall a waveform view which was stored with the provided sound files. These view memories are only used to speed up the tour; you could just as easily select each range by hand. You could also just as easily store your own waveform views in place of the ones which are there. The sound files supplied on the sound disk are normal Audio IFF format sampled sound files. When you've completed the Guided Tour, feel free to use them as such.

Now you've got all the information you'll need to get going, so set up your network and start the first mini-tour.

Setting Up Your Network

Alchemy is a program that can be personalized to reflect your own sampling network. The first time you use Alchemy, you should add all of your samplers to your program's network. Once you've done this a single time, Alchemy will remember what samplers you have, and where they are (in terms of MIDI channel, etc.). From then on, all of your samplers will appear as menu choices on Alchemy's Network menu. To set up your network, just make sure that you've plugged everything in as described in the Getting Started section of this manual. Then follow these directions:

Start up your Alchemy program by double-clicking the Alchemy icon.

When the Alchemy palette and desktop appear, you're ready to start.

Choose the New Instrument... command on the Network menu.

The New Instrument dialog box appears on your screen.

You'll complete this box once for every sampler you have in your network.

Use the pop-up Sampler ID menu at the top left of the box to select the name of the sampler you're adding.

Click on the button next to the type of communication you'll be using for this sampler.

This will probably be MIDI for most of your samplers, although a few machines can also use the higher-speed RS-422 or SCSI standards.

Click on the button next to the clock rate for the sampler you're adding.

Almost all MIDI samplers use a 1 Mhz clock rate.

Click on the port (either modem or printer) which is connected to the sampler you're adding.

If you're using a MIDI patcher with the sampler you're defining, then click on the port which is connected to the patcher.

If you're using a programmable MIDI patch bay, click in the box next to the words MIDI Patcher.

Alchemy can control your MIDI patcher to make sure that waveform information goes where it's supposed to.

If you're using a programmable MIDI patch bay, use the pop-up menus to select the patcher channel (channel that controls the patcher) and patcher program (for the specific sampler) which describe the sampler you're adding to the network.

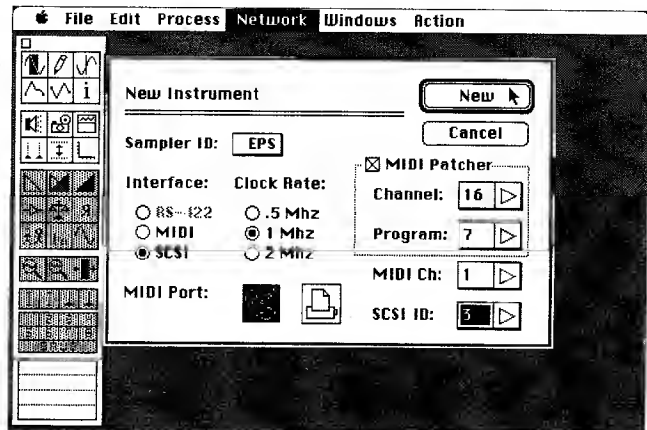
Use the pop-up menu to select the MIDI channel to which the sampler you're adding to the network is set.

If you are adding a SCSI sampler, use the pop-up menu to select the SCSI ID

to which the sampler you're adding to the network is set.

When you are all through, click on the New button at the bottom of the dialog box.

Your sampler has now been permanently added to the network. Check on the Network menu to make sure. The name of your sampler will now appear at the bottom. To add the rest of your samplers to the network, just select the New Instrument... command on the Network Menu and complete the dialog box one time for each sampler. Make sure to click on the New button after completing each dialog.



Note: If you've made a mistake, you can delete a sampler from the network by selecting the sampler's name at the bottom of the Network menu, and choosing the Delete Instrument... command again. When the dialog box appears for that sampler, just click on the Delete button. This removes the sampler from the network.

Your network is now set up, and you're ready to go.

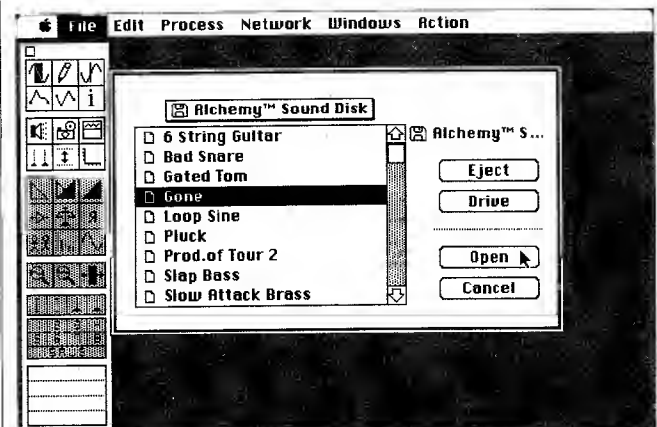
Mini-tour 1: Opening, Editing and Saving a Sound

This first mini-tour takes you through the simple process of opening, editing, and saving a sound on your Alchemy sound disk. For the purposes of this demonstration, you will find a sound file provided on your program disk. The file is called "Gone," and it contains a single spoken sentence. To progress through this mini-tour, start up the program and make sure that your network is set up according to the directions above. Then do the following:

Select the Open Special... command on the File menu.

When you do this the Alchemy Open Special... dialog box appears on your screen. If you are running the program from your hard disk, insert the Alchemy sound disk in any disk drive. This automatically displays the sound files on that disk. Notice how pertinent information is displayed about any selected sound file. If you are running Alchemy from a program diskette, insert the sound diskette in your other drive.

Click on the file named "Gone," and then on the Open button.



This opens a waveform window on your screen behind the Open Special... dialog box.

Click on the Quit button to close the Open Special... dialog.

The window containing the "Gone" sound appeared from behind the dialog. Notice that the waveform of the entire sound appears in the window.

Click on the Speaker icon on Alchemy's palette.



This plays back the entire sound that you just loaded, which consists of the sentence "I've really missed her since she's been gone." If you have a mini-plug hooked up to your Mac's audio out, the sound will play through that output (this is highly recommended). If nothing is plugged into the audio out, the sound will play through the Mac's internal speaker. *Note:* If you have a Sound Accelerator™ card installed, the sound will be played back at 16-bit resolution through the card's output jack. In all cases, playback volume is adjusted on the Macintosh's

Control Panel, and should be at "1" if you want line level audio out (see the menu).

Place the mouse cursor anywhere over the waveform, then click and drag it to select a waveform range.

By selecting a waveform range, you tell Alchemy what waveform section you would like to edit or hear. A selected range is always displayed in inverse color.

Click on the Speaker icon to hear the selected range.

When a waveform range is selected, it is the only thing that will play back when you click the Speaker icon (or press the space bar). You can now try selecting a number of different ranges

and listening to them. You'll find that it's quite easy to see each separate word in the sentence by looking at the contour of the waveform.

Choose the Select All command on the Edit Menu.

Your whole waveform is now selected and ready for editing. *Shortcut:* You can accomplish the same thing by double-clicking anywhere over the waveform.

Click on the Reverse icon on Alchemy's palette.



This reverses the selected range, which in this case is the entire waveform. The sound is now backwards.



Click on the Speaker icon to hear what you've done.

Click on the Reverse icon again to return the sound to normal.

You could have accomplished the same thing by selecting the Undo Reverse command on the Edit menu.

Click on the Axis Markers icon to display the rulers.



This helps you to measure the lengths of selected ranges, and will make it easier to work with the loop cursors by showing you a dark triangular tab under each. You can adjust the units displayed on the rulers by using the Axis Units pop-up menu on the Windows Menu.

Click on the Zoom box at the upper right corner of the waveform window's title bar

This grows your waveform window to a size that will make editing easier.

Click on the palette's Overview icon.



This creates an overview display above your waveform. The overview shows a map of your current waveform, and functions as a quick navigation tool.

Place the mouse cursor anywhere in the overview (upper) display and click-and-drag a dotted rectangle to choose a view range.

Notice how the waveform range you selected is automatically sized to appear in the waveform (lower) display, where you can edit it. That's how you use the overview display to navigate to any waveform range at any resolution. The white rectangle that appears over the overview display shows you what waveform range you're viewing below.

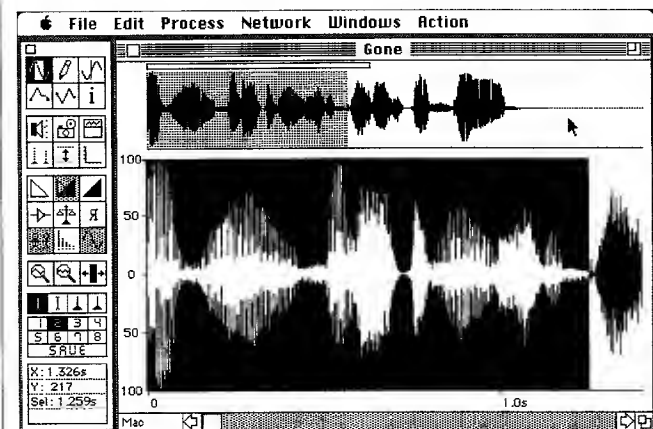
Click and drag to select a waveform range in the waveform (lower) display.

The gray gel in the overview display is updated to show the new selection range.

Click on the palette's view memory number 2.

1	2	3	4
5	6	7	8
SAVE			

Your waveform window is now zoomed in to show a large selected section at the beginning of the waveform. Notice how the overview display shows you where you are in the waveform. The gray gel (shaded area) in the overview display also shows you what waveform range is currently selected for editing.



Click on the Speaker icon to hear the selected waveform range.

Click on the palette's view memory number 1.

Now you're back where you started and ready to do some editing. Your aim will be to take the sentence "I've really missed her since she's been gone," and build the slightly more optimistic sentence "She's really missed me." Let's get started.

Click on view memory number 3.

The word "I've" is now selected. You can click on the Speaker icon (or press the space bar) to make sure.

Choose the Cut command on the Edit menu.

Notice how you just cut the work "I've" out of your waveform. The rest of the waveform slides over and no space is left.

Click on view memory number 4.

The word "She's" is now selected and ready for editing. Click on the Speaker icon to make sure.

Choose the Cut command on the Edit menu.

You've just cut the word "She's" out of your waveform.

Click on view memory number 1.

The blinking insertion point is now at the beginning of the waveform.

Choose the Insert command on the Edit menu.

A dialog box appears and asks if you want to increase your wavesample size.

Click on the Yes button.

You've just inserted the word "She's" at the beginning of your waveform.

Choose the Select All command on the Edit menu to select the whole waveform.

You can accomplish the same thing by double-clicking anywhere on the waveform.


Click on the Speaker icon to hear what you've got.

Click on view memory number 6.

The "m..." in "missed" is now selected for editing. Click the Speaker icon to make sure.

Choose the Copy command on the Edit menu.

A copy of the selected waveform range is now on the Mac Clipboard.



Click on view memory number 7, then choose the Paste command on the Edit menu.

You've now pasted the "m..." into place and you're ready for the "...eee" part.

Click on view memory number 5.

The "...eee" part of the word "She's" is now selected.

Choose the Copy command on the Edit menu.

A copy of your "...eee" is now on the Clipboard.

Click on view memory number 8.

The flashing insertion point is now located immediately after the "m..." which you pasted in earlier.

Choose the Insert command on the Edit menu.

A dialog box appears and asks whether to increase the wavesample's size.

Click on the Yes button.

You've just built the sentence "She's really missed me," but you're not done yet.

Double-click anywhere on the waveform, and then on the Speaker icon.

Notice there's still a bunch of leftover junk at the end of the waveform. You want to get rid of it.

Click on view memory number 2.

Hopefully the waveform range you want to keep is now selected. You can click on the Speaker icon to make sure. If the range is not exact, you can hold down the shift key and drag the selected range end so that only the "She's really missed me" sentence is selected.


Choose the Extract command on the Edit menu.

You've just used Alchemy to extract the good stuff and get rid of the bad. You now have what you set out to build. Although you have accomplished what you intended, there is one more tasteful touch you can add.

Choose the Auto Zero command on the Edit menu.

This assures you that any selected range is automatically adjusted to start and end on zero crossings. This may not mean much to you now, but with experience it will prove to be a useful tool in preventing clicks and pops in looping and scaling.

Click and drag the mouse cursor near the middle of the waveform until you have selected the word "really."



Be as accurate as possible. When you're satisfied with what you've got, continue.

Choose the Loop Selection command on the Edit menu.

This turns looping on, and automatically loops the word "really" which you had selected. Solid vertical loop cursors appear at the beginning and end of the range. Since you have the axis markers displayed, each loop cursor is marked by a black triangular base.

Double-click anywhere on the waveform to select the entire sound.

Click and hold down on the Speaker icon to hear the sound played up to its loop.

The loop repeats as long as you hold down the mouse button (or space bar). Click on the Loop Cursors icon if you want to turn off the loop again and hear the entire sound. Then click on it again to turn them back on. You're now ready to save what you've done.

Choose the Save As... command on the File menu.

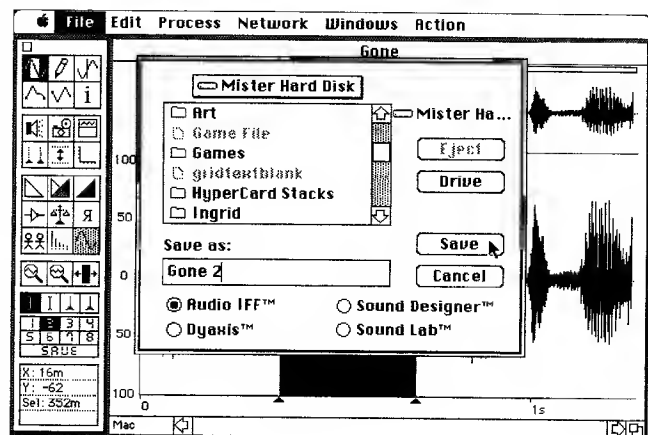
When you do this, a dialog box appears on your screen. Use this box just like a standard Macintosh Save As... dialog box to select the disk and/or folder where you want to store your new file. Then type in the new file name. When you have chosen a destination and named your new file, continue.

Click on the Audio IFF button to choose your file format.

The Audio IFF format is the default format, and is the most versatile (IFF stands for Interchange File Format). It can be mono or stereo, and it is the only format that stores view memories.

Click the Save button to actually save it.

You have just successfully opened an old sound file, edited it extensively, and saved it as a new file. The new sound file you created can be opened at any time, and sent to any sampler in your network. Obviously this mini-tour was partially automated (using the view memories), and you really only used rudimentary tools. You can be sure that you haven't stretched Alchemy anywhere near its limits. In the next mini-tours you will have a bit more freedom to do so.



Mini-tour 2: Creating, Looping, and Transferring a Sound

This mini-tour takes you through the simple steps required to open multiple sound files, create and loop a totally new sound, and send that sound to a sampler. Alchemy makes this type of sound design easy by allowing you unlimited open windows, and central transfer to any network sampler. To begin this mini-tour, make sure that you've set up your network as described near the beginning of this chapter. Then start up Alchemy by double-clicking on its icon. When the Alchemy palette and desktop appear, you're ready to begin.

Important Note: If your Macintosh doesn't have enough available memory to open all of the sound files in this mini-tour simultaneously, open them one at a time when editing is required.

Choose the Open Special... command on the File menu.

When you do this the Alchemy Open Special dialog box appears on your screen. If you are running the program from your hard disk, insert the Alchemy sound disk in any disk drive. This automatically displays the sound files on that disk. If you are running Alchemy from a program diskette, insert the sound diskette in your other drive. The Open Special... command will keep appearing until you click on the box's Quit button.

Click on different files, and then on the Listen button.

Selecting any sound file shows you all of its pertinent information. The Listen button plays the sound directly from hard disk. The listen button can be used to

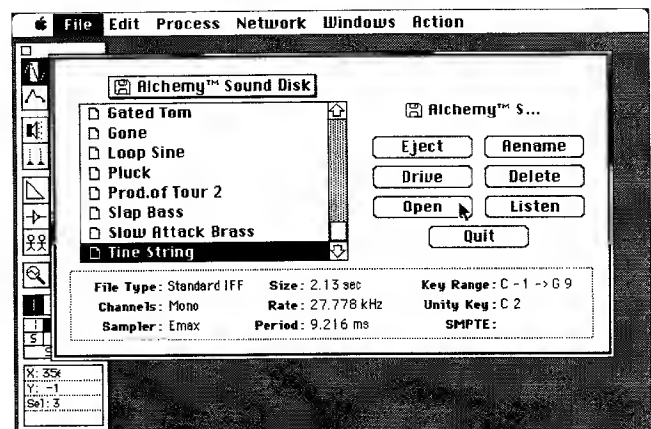
audition sounds which are larger than memory, and they will be played back in 16-bit stereo if you have a Sound Accelerator™ card installed.

Click on the file named "Tine String," and then on the Open button.

This opens a waveform window on your screen which contains the "Tine String" sound. The waveform of the entire sound appears in the window.

Repeat the above procedure to open the files entitled "Slow Attack Brass" and "6 String Guitar."

Each new window appears almost on top of the window before it. You'll be fixing this in a second.



Click on the open dialog box's Quit button.

This closes the Open Special dialog.

Select the New command on the File menu.

A fourth window appears on top of the others, but this one is empty. This will be your waveform work area.

Choose the Tile command on the Windows menu.

The Tile command takes all of the windows which are presently open and places them in even squares across the screen.

If the axis markers are not showing, click on the Axis Markers icon.



The axis markers generally make working with multiple sounds much easier, because they show all relative lengths and amplitudes.

Activate the "Tine String" window.

You do this by clicking the mouse cursor anywhere in the "Tine String" window. The window is now ready for editing and playback.

Click on the Speaker icon to hear the sound.

If you have a mini-plug hooked up to your Mac's audio out, the "Tine String" sound will play through that output (or through a Sound Accelerator™ card, if installed). If no card is present, and nothing is plugged into the audio out, the sound will play through the Mac's internal speaker.


Note: Playback volume is adjusted on the Macintosh's Control Panel (see the menu). A volume of "1" is recommended for line level output.

Use the mouse to select each window one at a time and click on the Speaker icon (or press the space bar) to hear the sounds.

Notice how looped sounds are played back indefinitely until you let go of the mouse button. The sounds which are now on your screen are the building blocks of the sound you're about to construct. Your goal is to create a 1.9 second sound with a "Tine String" beginning, a "Brass" ending, and a "6 String Guitar" sound strongly mixed in for the duration. Let's get started.

Click the mouse on the "Tine String" window to select it. Then click in the little "Zoom box" located at the top right corner of the window's title bar.

Clicking in a window's zoom box always grows the window to cover the screen. Clicking in the zoom box a second time returns it to its



previous size. It's usually easier to accomplish precise edits when a waveform is displayed as large as possible.

Choose the Select All command on the Edit menu.

The whole waveform is now selected and ready for editing. Take a look at the numeric display at the bottom of Alchemy's palette. Notice that the length of the selected waveform (Sel:) is 2.129 seconds. Since you want to construct a sound which is 1.9 seconds long, you'll want to select the first 1.9 seconds of the waveform.

Click on view memory number 2.

If you look in the palette's numeric display again, you'll find that a 1.9 second range is now selected. Although you could have easily done this by hand, view memories have been included to help speed up the process.

Choose the Copy command from the Edit menu.

A copy of the 1.9 second "Tine String" range is now on the Mac Clipboard where it will wait to be pasted, mixed, or inserted.

Click on the "Tine String" window's zoom box again.

This shrinks it down to its previous size.

Click the mouse on the "Untitled" waveform window to select it. Then choose the Paste command on the Edit menu.

A dialog box appears and asks if the wavesample size should be increased.

Click on the Yes button.

A copy of the 1.9 second "Tine String" sound is now in the "Untitled" window, and you have now accomplished your first step.

To zoom all the way out and see your whole waveform, choose the Full Zoom

Out command on the Action menu.

You can accomplish the same thing by holding down the command (⌘) key and clicking on the Zoom Out icon.

Click on the Fade Out icon.




The "Tine String" sound in the Untitled waveform window is now faded out.

Click on the Speaker icon to hear the copy.

Click on the palette's SAVE button, and then on view memory number 1.

You've just saved your current waveform view as view number 1. From now



on, you can recall this view by selecting the “Untitled” window and clicking on view memory number 1. Each window can store eight separate views permanently.

Click on the “Slow Attack Brass” window to select it.

You can click on the Speaker icon to remind yourself how it sounds.

Click on view memory number 2.

If you look in the palette’s numeric display, you’ll see that a waveform range of 1.9 seconds is now selected and ready for editing.

Click on the “Slow Attack Brass” Zoom box.

This grows the waveform window to cover the screen so it’s easier to edit.

Click on the Threshold Bars icon to show the threshold bars in the active window.



The threshold bars appear in the “Slow Attack Brass” window. The threshold bars are dotted horizontal lines at the top and bottom of the

window. You’ll use these bars to set a scaling factor for amplitude change. You’re about to scale down the “Slow Attack Brass” sound to reduce its amplitude and prepare it for mixing.

Place the mouse cursor over the upper threshold bar in the “Slow Attack Brass” window.

The cursor changes to an up and down arrow.

Click and drag the bar downwards to scale the selection to 70%.

Look at the palette’s numeric display as you do this. Keep dragging the threshold bar downwards until the Scale: box reads “0.70.” When you eventually execute the amplitude change, this


setting will decrease the amplitude of the selected waveform range to 70% of its former level.

Click on the Scale icon to execute the amplitude change.

You can click on the Speaker icon to hear how this has affected the sound, but effects should also be noticeable in the visual display. The scaling process only takes place over the selected waveform range.

Click on the Fade In icon to taper the amplitude of the selected “Slow Attack Brass” range.

Click on the Speaker icon to hear what you’ve done. Notice once again that only the selected range is effected by the fade.



Choose the copy command from the Edit menu.

A copy of the scaled and faded in “Slow Attack Brass” waveform range is now in the Mac Clipboard.

Click on the “Slow Attack Brass” window’s zoom box.

This shrinks the waveform window back down to its previous size.

Click on the “Untitled” window to select it for editing.

Choose the Mix command on the Edit menu.

The faded in “Slow Attack Brass” waveform has now been mixed with the faded out “Tine String” waveform. Click on the Speaker icon to hear the result. The transition from Tine to Brass should be smooth and natural.

Click on the Threshold Bars icon to show the threshold bars in the “Untitled” window.

The threshold bars appear as before. Once again, you’re going to scale down a sound so it won’t be as prominent when you mix it.

Click on the “Untitled” window’s zoom box.

This grows the waveform window to cover the screen so it’s easier to edit.

Place the mouse cursor over the upper threshold bar and click and drag the bar downwards to set a scaling factor of 80%.

Look at the palette’s numeric display as you do this. Drag the threshold bar downwards until the Scale: box reads “0.80.” When you eventually click on the Scale icon, this setting will decrease the amplitude of the selected waveform range to 80% of its former level.

Click on the Scale icon to execute the scale.



You can click on the Speaker icon to hear how this has affected the sound.

Click on the “Untitled” window’s Zoom box.

This shrinks the window back down to its previous size.

Click on the “6 String Guitar” window to select it for editing.

Click on view memory number 2.

If you look at the palette’s numeric display, you’ll find that a 1.9 second section of the waveform is now selected.

Choose the Copy command on the Edit menu.

A copy of the selected “Guitar” waveform range is now on the Mac Clipboard and is ready for pasting, mixing or inserting.

Click on the “Untitled” waveform window to select it for editing.

Choose the Mix command on the Edit menu.

The “Guitar” waveform from the Clipboard is now mixed with the Tine/Brass waveform which you constructed earlier.

Click on the Speaker icon to hear what you’ve built.

You’ve now just about completed your task. You have a sound with a “Tine String” beginning, a “Brass”

ending, and a strong “6 String Guitar” mixed in for its duration. To really complete the sound, you’ll need to loop it.

Click on the “Untitled” window’s Zoom box to enlarge it for easy editing.

Click on the Threshold Bars icon to turn off the threshold bars.

Choose the Auto Zero command on the Edit menu.

When Auto Zero is set, any waveform range you select will be automatically adjusted to begin and end on zero crossings. This will help prevent accidental clicks and pops in the final waveform.

Use the mouse to select approximately the last 1/4 of the waveform.

Try not to select any more than the last 1/4 or you’ll get more harmonic evolution than you’d like. This will be your basic loop area.

Choose the Loop Selection command on the Edit menu.

This automatically turns the loop cursors on and loops the selected range.

Click and hold on the Speaker icon .

Unless you’re incredibly lucky, your will have a click, pop, or bump in it. Don’t worry, there’s a quick way to fix this.

Choose the Xfade Loop... command on the Process menu.

When the crossfade looping dialog box appears, click on the OK button.

Alchemy automatically optimized your loop by executing the crossfade over approximately 1/4 of your loop range. This can be adjusted if you wish to have a larger or smaller crossfade range.

Click and hold on the Speaker icon again.

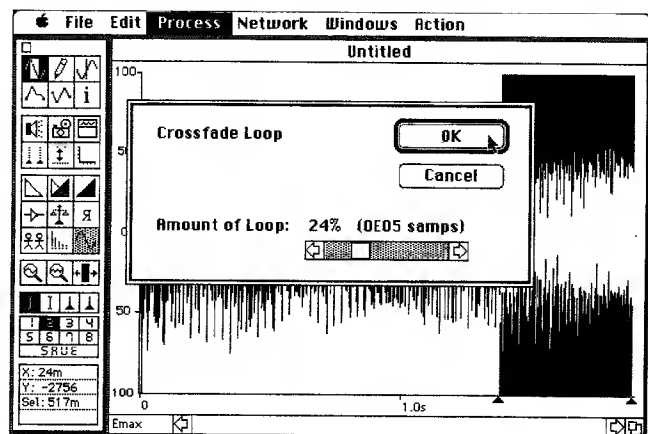
The crossfade loop should sound quite a bit better than your original test loop.

Double-click anywhere on the untitled waveform to select the entire sound.

You can click and hold on the Speaker icon to hear your new product. You’ve now completed your sound and you’re ready to save it.

Choose the Save As... command on the File menu.

When the Save As... dialog box appears, make sure the Audio IFF button has a black dot inside it. If it doesn’t, click on it. Then use the dialog box like a standard



Macintosh Save As... box to select the disk and/or folder where you want to store the new file, and type in the new file name.

When you've filled out the Save As... dialog box, click on the Save button to finalize your save.

Now you're ready for the final step: transferring your new sound to a sampler. Before you do this, make sure that you followed the directions in the Setting Up Your Network section above, and that your sampler is on line and booted up with a disk.

Choose the name of the destination sampler at the bottom of the Network menu.

When the sampler is selected, a check mark appears in front of its name on the menu.

Chose the Send Sound command on the Network menu.

The keyboard dialog box appears.

If it is not already checked, click in the Assign New Voice box.

The entire keyboard is automatically highlighted as the destination range, because that is the default key range for the untitled file you created with the New...command. A sound file's key range is remembered when it is saved, and the Assign New Voice box always defaults to the sound file's key range. A sound file's native key range may be viewed and edited using the Soundfile Setup... command on the File menu.

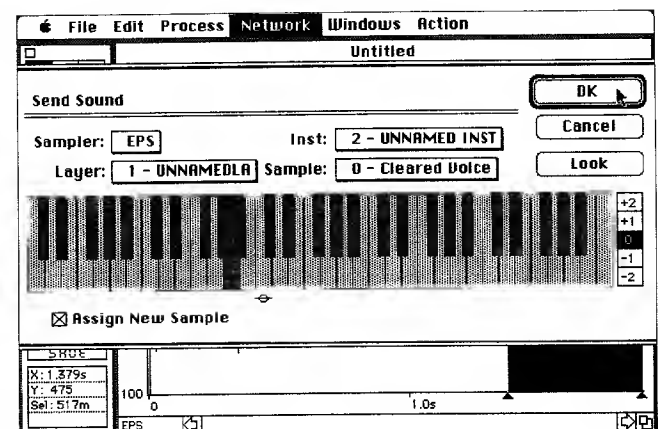
Click on the single dark gray key and drag it to a position towards the middle of the keyboard.

The dark gray key is the Unity key, and represents the actual unity playback key which will play the sound

without transposition upwards or downwards.

When you're finished, click OK.

If you have the Alchemy's auto resample preference chosen, you will be prompted to choose a valid sample rate, and verify the adjustment to overall sound length. If you are assigning the new sound over existing sounds, a number of dialog boxes may appear asking if you want to delete voices already in memory. Don't worry, this doesn't delete them from your disk, it just replaces them in your sampler's memory. Click OK in each box that appears. When you've done this, a thermometer-style bar appears to show you how the transfer is going.



When the transfer is complete, play your sampler to hear your sound.

It may not be the best sound you've ever heard, but that's not your fault. At the very least it taught you something about Alchemy, and that's what this tour is all about.

You have now created, edited, saved, and transferred a new sound. Once saved, the sound file, along with its loop and key range information, will always remain on disk for later editing. After you have transferred a sound to your sampler, you may want to save it on a performance diskette for loading in live environments. If you use your samplers primarily in the studio, make sure you leave a copy of the sound on your Macintosh hard disk as the beginning of a universal sample library.

Mini-tour 3: Creating a Stereo Sound File

Alchemy's sound design environment allows you to make and keep both mono and stereo files. Stereo sound files may consist of actual stereo sampled sounds (from an Emulator III or Dyaxis, for example), of matched mono samples of a true stereo image, or of hand-built stereo images. Once you've captured or created a stereo sound file, Alchemy takes care of transferring, assigning, and panning the left and right waveforms correctly. To get a good idea of how this works, make sure that your network is set up correctly (see Setting Up Your Network at the beginning of this section) and go through this mini tour.

Important Note: Stereo sound files may only be played back in stereo by samplers which have two or more assignable outputs, and which can trigger at least two sounds from the same key range. Stereo sounds may also be played back directly from the Mac at true 16-bit

fidelity, if you have the Sound Accelerator™ card.

Select the Open... command on the File menu.

When you do this the standard Macintosh Open dialog box appears on your screen. If you are running the program from your hard disk, insert the Alchemy sound disk in any disk drive. This automatically displays the sound files on that disk. If you are running Alchemy from a program diskette, insert the sound diskette in your other drive.

Click on the file named "Slap Bass" and then on the Open button.

This opens a waveform window on your screen which contains the "Slap Bass" sound. The waveform of the entire sound appears in the window.

Click on the Speaker icon to hear the sound.

Choose the Mono to Stereo command on the File menu.

Your original file is now a stereo file containing two waveforms. The upper is the left channel and the lower is

the right channel. The waveforms in both channels are now identical, but you'll be changing that.

Double-click the mouse cursor anywhere in the lower (right channel) waveform.

The whole right channel is now selected for editing.

Click on the Reverse icon .

You now have stereo image which begins with a slap bass attack in the left channel and ends with a backwards slap bass crescendo in the right channel. The image would give the impression of a movement from left to right.

Important Note: If you have a Mac II or SE and Sound Accelerator™ card, you will be able to click on the Speaker icon and hear both channels of the stereo image. With any other Mac, only one channel may be played back at a time. If two channels are selected, only the left one plays. To really hear the image, you'll need to transfer the sound to a stereo playback sampler.

Click on the Speaker icon to hear what you have.

If you don't have a capable Mac, you'll need to select and listen to each channel separately.

Click on the Axis Markers icon.

The axis markers are now showing.

Select the Open... command on the File menu.

The standard Macintosh Open dialog box appears on your screen.

Click on the file named "Pluck" and then on the Open button.

This opens a waveform window that contains the "Pluck" sound. The waveform of the entire sound appears in the window.

Choose the Strip command on the Windows menu.

This cleans up the display by placing all open sound files in horizontal windows of equal size.


Double-click anywhere on the "Pluck" waveform to select the entire sound.

Click on the Speaker icon to hear the sound.

You will be adding this sound more or less randomly into the "Slap Bass" stereo image.

Choose the Copy command on the Edit menu.

A copy of the "Pluck" waveform is now on the Mac Clipboard.



Activate the stereo “slap bass” window by clicking it.

Place the mouse cursor over the right (lower) channel.

Look in the palette’s numeric display next to “X:” to see what fraction of a second the cursor is over.

Click the mouse a single time on 115m (milliseconds).

The flashing insertion point now indicates your insertion point.

Choose the Mix command on the Edit menu.

A copy of the “Pluck” waveform is now mixed in near the beginning of the right channel waveform.

Double-click on the lower waveform to select the entire channel.

Click the Speaker icon or press the space bar to hear what you’ve added.

Place the mouse cursor over the left (upper) channel.

Look in the palette’s numeric display next to “X:” to see what fraction of a second the cursor is over.

Click the mouse a single time on approximately 550m (milliseconds).

The flashing insertion point now indicates your insertion point.

Choose the Mix command on the Edit menu.

A copy of the “Pluck” waveform is now mixed in near the end of the left channel waveform.

Choose the Select All command on the Edit menu to select the entire channel.

Click the Speaker icon or press the space bar to hear what you’ve added.

Place the mouse cursor over the channel separator dividing the right channel from the left.

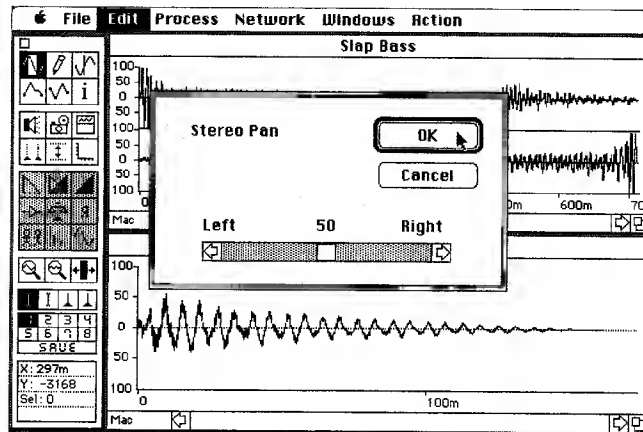
This is how you select a waveform range in both channels at once. Look in the palette’s numeric display next to “X:” to see what fraction of a second the cursor is over.

Click the mouse a single time on 320m (milliseconds).

The flashing insertion point now indicates your insertion point.

Choose the Mix command on the Edit menu.

A dialog box appears and asks you how you’d like to pan the Clipboard waveform into the stereo image.



Click on the OK button.

A copy of the "Pluck" waveform is now mixed in and panned to the center of both the left and right channel waveforms.

Double-click anywhere over the channel separator to select both channels of the entire waveform.

Both channels are now selected.

Choose the Loop Selection command on the Edit menu.

You've just looped the entire waveform for playback. If your Mac is capable of stereo playback, click on the Speaker icon to hear it. Otherwise you'll be saving it and transferring it to a sampler to listen to it.

Choose the Save As... command on the File menu.

When the Save As... dialog box appears, make sure the Audio IFF button has a black dot inside it. If it doesn't, click on it. Then use the dialog box like a standard Macintosh Save As... box to select the disk and/or folder where you want to store the new file, and type in the new file name.

When you've filled out the Save As.. dialog box, click on the Save button to finalize your save.

Now you're ready for the final step: transferring your new sound to a sampler. Before you do this, make sure that you followed the directions in the Setting Up Your Network section above, and that your sampler is on line and has its system software loaded.

Choose the name of the destination sampler at the bottom of the Network menu.


When the sampler is selected, a check mark appears in front of its name on the menu.

Choose the Send Sound command on the Network menu.

The keyboard dialog box appears.

Click in the Assign New Voice box.

The entire keyboard is automatically highlighted as the destination range, because that is the default key range for the untitled file you created with the New... command. A sound file's key range is remembered when it is saved, and the Assign New Voice box always defaults to the sound file's key range. A sound file's native key range may be viewed and edited using the Soundfile Setup... command on the File menu.



Click on the single dark gray key and drag it to a position towards the middle of the keyboard.

The dark gray key is the Unity key, and represents the actual unity playback key which will play the sound without transposition upwards or downwards.

When you're finished, click OK.

A number of dialog boxes may appear asking if you want to delete voices already in memory. Don't worry, this doesn't delete them from your disk, it just replaces them in your sampler's memory. Click OK in each box that appears. When you've done this, a thermometer-style bar appears to show you how the transfer is going. One transfer thermometer appears for each of the two channels.

When the transfer is complete, play your sampler (preferably near the middle key ranges) to hear your sound.

Your stereo sound is now assigned over the entire keyboard. If you have your sampler's stereo outputs hooked up, you'll hear the

original forwards/backwards "Slap Bass" undulating from far left to far right in the stereo image. At the same time, the three "Pluck" sounds you mixed in are crossing the slap bass sound and swinging from right to left in the stereo image. If this is not what you hear, make sure that you have your sampler hooked up to play back in stereo.

That concludes your stereo mini-tour. Although you've hand-built a stereo effect, you'll find that it's also possible to put left and right channels of an ambient stereo image in a stereo sound file. If you'd like to learn more about how your specific sampling devices achieve stereo, see Using Alchemy, and the sampler-specific information in the Appendix of this manual.

Mini-tour 4: Editing a Sound's Harmonic Spectrum

The fourth mini-tour offers you a closer look at some of Alchemy's more powerful digital signal processing features. In it you will have a chance to use Alchemy's harmonic spectrum display to analyze and edit the harmonic spectrum of a sound. This type of editing is possible thanks to the Fast Fourier Transform.

The Fast Fourier Transform (FFT) is a powerful algorithm which you can use to determine exactly which frequencies are present in any waveform, and their precise amplitudes. Once you have this frequency information, you are free to change that harmonic spectrum in any way and then reverse the transform. This changes the wave shape in an extremely precise way. You could not accomplish the same thing by simply redrawing the waveform by hand.

If you'd like some more background information

about this, take a look at the About Sound section of this manual. If you're ready to try it out, continue on.

Select the Open... command on the File menu.

When you do this the standard Macintosh open dialog box appears on your screen. If you are running the program from your hard disk, insert the Alchemy sound disk in any disk drive. This automatically displays the sound files on that disk. If you are running Alchemy from a program diskette, insert the sound diskette in your other drive.

Click on the file named "Gated Tom," and then on the Open button.

This opens a waveform window on your screen which contains the "Gated Tom" sound. The waveform of the entire sound appears in the window.

Click on the Speaker icon on Alchemy's palette.

This plays back the entire sound that you just loaded, which is a basic gated tom sound. If you have a mini-plug hooked up to your Mac's audio out, the sound will play

through that output (or through the Sound Accelerator™ card, if you have one)). If no card is installed, and nothing is plugged into the audio out, the sound will play through the Mac's internal speaker. *Note:* Playback volume is adjusted on the Macintosh's Control Panel (see the menu). A volume of "1" is suggested for line level output.

Click on the waveform window's Zoom box.

The Zoom box is located in the right corner of the window's title bar. It grows the window to the largest size possible, which makes the sound easier to see and edit. Clicking on the Zoom box a second time returned the window to its previous size.


Click on the palette's Overview icon.



An overview is created above the current waveform. This will help you keep track of where you are in the "Gated Tom" sound.

Click on view memory number 2.

You are now zoomed in and looking at a single waveform period near the beginning of the sound. If you're unclear what is meant by "period," you can find out in the About Sound section of this manual. Notice how the overview display shows you both what you're viewing in the lower display (the white rectangle) and what's currently selected (the gray gel).



Click on the Analyze icon.



The single selected waveform period has now undergone Fourier analysis, and a harmonic spectrum window has been opened to describe it. This window details the spectral content of the selected waveform range. Each vertical bar represents a frequency, and its height tells you how much of that frequency component is present. Notice that the tallest bar (channel) is located at the far left of the display. This is the strongest frequency component in the analyzed sound, and is called the "fundamental."

Click the mouse cursor above the highest vertical bar in the newly opened harmonic spectrum display.

A solid black square appears at the top of the bar. Now check in the palette's numeric display. The frequency of the selected bar is shown in the second box from the top. You'll notice that the fundamental frequency of the analyzed waveform range is 125.7.

Click drag the fundamental downwards until it has no amplitude.

You can adjust the level of any frequency component by clicking and holding on the selection box at its tip and dragging the channel to a new amplitude. You'll know that your component has zero amplitude when the "dB:" box in the numeric display contains the "--" symbols. You can also delete a selected frequency by choosing the Cut command on the Edit menu.

Click on the "Gated Tom" waveform window to make it active.

This will allow you to watch as the new waveform is resynthesized. The harmonic spectrum window is still open; it's just behind the waveform window.

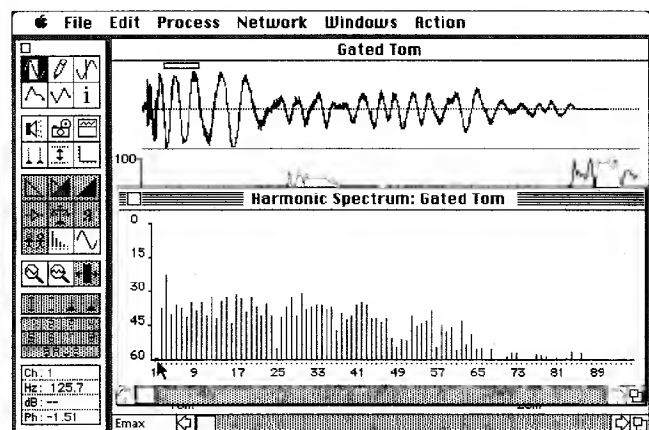
Click on the Resynthesize icon.



You have now resynthesized your originally analyzed waveform to reflect your deletion of the fundamental frequency. You can tell from the reduced amplitude of the new waveform that a very major harmonic component has been removed.

Choose the Undo command on the Edit menu.

Now you can compare the changed waveform to the original. If you want to compare them again, just choose the Redo command on the Edit menu, and then the Undo command again.



Hold down the ⌘ key and click on the Zoom Out icon.

The waveform view is zoomed all the way out to show the whole sound. You can also accomplish this by choosing the Full Zoom Out command on the Action menu.

Double-click anywhere on the waveform to select the whole sound.

Click on the Analyze icon.

This time you have selected the entire sound for spectral analysis, which means you'll be able to edit the spectrum of the whole waveform. More data is being analyzed, so it may take a moment before the harmonic spectrum window appears.

Select a single frequency channel of approximately 450 Hz.

You can tell the frequency of any selected channel by looking in the numeric display. 450 Hz. should be somewhere around the 90th frequency channel.

Choose the Clear Below command on the Edit menu.

This command is only enabled when you're working in a harmonic spectrum window. When you choose it you see all of the frequency bands below the selected channel removed from the spectrum. The Clear Below command only functions when a single frequency channel is selected.

Click on the "Gated Tom" waveform window to make it active.

This will allow you to watch as the new waveform is resynthesized. The harmonic spectrum window is still open; it's just behind the waveform window.

Click on the Resynthesize icon.

You have now resynthesized the entire waveform to reflect your removal of all frequencies below 450 Hz. You can tell from the shape of the new waveform that a very major change has taken place.


Click on the Speaker icon to hear the new waveform.

Notice how you can both see and hear the total lack of low end frequencies. You've constructed a very accurate high-pass filter.

Choose the Undo command on the Edit menu.

Now you can compare the changed waveform to the original.

Click on the Speaker icon to hear the original waveform.



There's obviously quite a difference between the two waveforms. If you want to compare them again, just click on the Redo command and then the Speaker icon to hear the altered waveform. The Undo command will bring back the original.

This illustrates some of the power of harmonic spectrum editing. By using the harmonic spectrum window, you were able to take an existing sampled waveform and edit its frequency content directly. The FFT algorithm took care of rebuilding the new waveform to reflect the spectral changes you made.

FFT analysis and resynthesis can be used as a tool to determine the exact spectrum of a waveform, and it can be used to change that waveform in a very frequency-accurate fashion. For isolation of frequency content, and construction of razor sharp filters (like your high-pass filter) it is indispensable.

Mini-tour 5: Using Alchemy's Enveloping, Time Scaling, and Pitch Shifting

The fifth mini-tour introduces you to some of the powerful new features that have been added to Alchemy 2.0. This newest version of Alchemy contains two entirely new editing modes, and a number of new DSP functions that greatly increase the scope of the program. In this mini-tour you will use both amplitude and frequency enveloping to create a sound. Then you will adjust the pitch of that sound without adjusting its duration.

Amplitude enveloping is the process of superimposing the envelope of one sound over another. You might, for example, trace the envelope of a violin, and superimpose that envelope over a piano to alter the way its volume changes over time. And that is an example of only the simplest level of amplitude enveloping capabilities.

Frequency enveloping is the process of modulating the

frequency of one sound by another sound. At its most basic level, you can use this function to design totally new sounds that are the product of FM synthesis using complex sampled waveforms.

Add to this Alchemy 2.0's new Time Scale and Pitch Shift functions, and you have a relatively limitless collection of sound design tools. Time scaling is the process of changing the duration of a sampled sound without altering its pitch. This offers a long-awaited capability to those involved in audio post-production, because it facilitates the adjustment of sounds to match particular time (or SMPTE) durations. Pitch shifting allows you to change the pitch of a sampled sound accurately, with *or without* changing its duration. This makes it possible for you to build a number of different effects, including complex harmonizing.

Alchemy 2.0's new modes and functions are very sophisticated, and to truly understand them you should read through the Using Alchemy chapter of the manual. But, to get a taste of their capabilities, proceed with this mini-tour.

Note: Bear in mind that it may take up to five minutes to pitch shift the tour sound, if you do not have the Sound Accelerator™ installed.

Select the Open... command on the File menu.

When you do this the Alchemy open dialog box appears on your screen. If you are running the program from your hard disk, insert the Alchemy sound disk in any disk drive. This automatically displays the sound files on that disk. If you are running Alchemy from a program diskette, insert the sound diskette in your other drive.

Click on the file named “6 String Guitar,” and then on the Open button.

This opens a waveform window on your screen which contains the “6 String Guitar” sound. The waveform of the entire sound appears in the window.

Click on the Speaker icon on Alchemy’s palette to hear the sound.

Your first step will be the construction of a harmonizing effect.

Click on the Loop Cursors icon to turn off the loop.

Double-click anywhere on the waveform to select the entire sound.

The whole waveform is now ready for editing.

Choose the Pitch Shift command on the Process menu.

This opens the pitch shift dialog, which you can use to adjust the pitch of any sound, with or without changing its duration.

Type “-5” into the “Transpose by” box.

You will be lowering the pitch of the selected range by five semi-tones.

Click on the box in front of Preserve Duration.

This is a sophisticated function that will make sure that the duration of the sound is not changed by the pitch shift.


Click on the OK button.

The “6 String Guitar” window will reappear containing the pitch shifted waveform.

Click on the Speaker icon on Alchemy’s palette to hear the pitch shifted sound.

Choose the Copy command on the Edit menu.

A copy of the affected waveform is now on the Clipboard.



Choose the Undo command on the Edit menu.

The “6 String Guitar” sound is now restored to its original pitch.

Choose the Mix command on the Edit menu.

In effect, you have just harmonized your original sound by mixing it evenly with a lower-pitched version of itself.

Click on the Speaker icon on Alchemy’s palette to hear the effect you’ve built.

Select the Open... command on the File menu.

The Alchemy open dialog appears.

Click on the file named “Loop Sine,” and then on the Open button.

The “Loop Sine” waveform appears in its own window for editing.

Double-click anywhere on the “Loop Sine” waveform to select the entire sound.

Choose the Copy command on the Edit menu.

A copy of the “Loop Sine” waveform is now on the Clipboard.

Close the Loop Sine window by clicking on the Close Box in its title bar.

The “Loop Sine” window disappears from your screen.

If it is not in front, activate the “6 String Guitar” window by clicking on it.

The active window is always indicated by horizontal “grabber bars” in the title bar.

Switch to Frequency Enveloping mode by clicking on the Frequency Modulation mode icon.




The waveform will now appear gray, and cannot be edited. The black horizontal line on the X axis indicates the current frequency modulation envelope, which is set flat at zero (no modulation). This envelope can be dragged to any new profile, or drawn using the pencil tool (see Using Alchemy for more information).

Use the Frequency Range pop-up menu on the Process menu to select “1 Octave” as the modulation range.

Make sure that the axis units are showing. If they are not, click on the Axis Markers icon to display them.

Choose the Paste Envelope command on the Edit menu.



You have now pasted the Clipboard “Loop Sine” waveform into this frequency modulation window as the modulation envelope. Alchemy treats all sampled waveforms the same. Any envelope can be used as a waveform, and any waveform can be used as an envelope.

Click on the palette’s pitch shift icon.



The “6 String Guitar” waveform is now being frequency modulated according to the “Loop Sine” modulation envelope. The pitch is being shifted smoothly between one octave above and one octave below normal pitch.

Click on the Speaker icon on Alchemy’s palette to hear the Frequency modulation.

The undulating pitch modulation will be evident. You have finished your pitch modulation effect.

Select the Open... command on the File menu.

The Alchemy open dialog appears.

Click on the file named “Pluck,” and then on the Open button.

The “Pluck” waveform appears in an open window, ready for editing. You will be isolating the amplitude envelope of this waveform and superimposing it over the product of your “6 String Guitar” edits.

Double-click anywhere on the “Pluck” waveform to select the entire sound.

Switch to Amplitude Modulation mode by clicking on the Amplitude Modulation mode icon.




The waveform will now appear gray, and cannot be edited. The black horizontal

line near the top of the Y axis indicates the current amplitude modulation envelope, which is set to full on (all attack). Like the frequency modulation envelope, this envelope can be dragged to any new profile, or drawn using the pencil tool (see Using Alchemy for more information).

Choose the Trace Envelope command on the Process menu.

After a moment the amplitude envelope of the “Pluck” sound will appear as a black line in the window. You will notice that it outlines the overall amplitude of the “Pluck” waveform. Unlike frequency envelopes, amplitude envelopes are always positive.



Choose the Copy Envelope command on the Edit menu.

A copy of the “Pluck” amplitude envelope is now on the Clipboard.

Activate the “6 String Guitar” window by clicking on it.

The active window is always indicated by horizontal “grabber bars” in the title bar.

Switch to Amplitude Modulation mode by clicking on the Amplitude Modulation mode icon.

The “6 String Guitar” waveform will now appear gray, and cannot be edited.

Choose the Paste Envelope command on the Edit menu.

The amplitude envelope you originally traced and copied from the “Pluck” waveform is now pasted into the “6 String Guitar” amplitude window.

Click on the Scale icon to scale the amplitude envelope to maximum amplitude value.

Most of the normal editing functions (reverse, invert, scale, face out, and fade in) can be executed on envelopes as well as waveforms. By scaling up the envelope, you can be assured that the resulting waveform will make use of the full amplitude scale.

Click on the Amplitude fit icon to fit the “6 String Guitar” waveform to the “Pluck” envelope.



The normal amplitude of the “6 String Guitar” sound is now adjusted to fit exactly the “Pluck” amplitude envelope. Although the tone of the “6 String Guitar” sound remains unchanged, its amplitude now develops over time just as the “Pluck” sound develops.

Switch back to Range Selection mode by clicking on the Range Selection mode icon.

The “6 String Guitar” waveform now appears in black, and is ready for editing.

Click on the Speaker icon to hear the product of your work.

By building a pitch-shift harmonizer, and then performing both frequency and amplitude enveloping, you have created an entirely new sound, with some of the tonal qualities of the “6 String Guitar” sound, and the amplitude profile of the “Pluck” sound. Although this is a very rudimentary use of Alchemy 2.0’s time scaling, pitch shifting, and enveloping functions, perhaps it gives you an idea of what these functions can do.

With the completion of this final mini-tour, you have successfully tried many of Alchemy’s time domain editing, frequency domain editing, file manipulation, enveloping, and network transferring functions. To balance out your knowledge and fill in the blanks, read through the Using Alchemy section which comes next. It will be your key to getting the most out of this program and your network.

Using Alchemy

Chapter Three

Introduction

By now you should have a basic feel for some of the things that Alchemy can do. If you've gone through the Guided Tour section of this manual, you've been introduced to some of Alchemy's waveform and harmonic editing capabilities, and you've seen how Alchemy can be used to share sampled sounds over a Distributed Audio Network (DAN).

This section of the manual is where you'll really learn what Alchemy is all about. Alchemy is a universal stereo sample editing network, and because of that it's revolutionary in many ways. To truly unleash Alchemy's power, spend some time learning the different tools available to you. Learning to edit waveforms in the time domain is a skill in itself, as is editing a waveform's harmonic spectrum. To really grasp how all of this works, go through this section sequentially, and try the different tools and processes yourself as you read about them.

Once you have gone through this Using Alchemy section, you'll have all of the information you need to put Alchemy through its paces. Alchemy is very open-ended, and the way you use it will depend on what you're using it for. Whether you're editing stereo sound effects, or just looping a sampled voice for your keyboard, experience will be your best teacher.

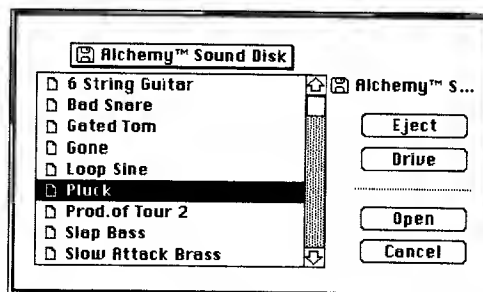
Opening and Saving Files

Alchemy has been designed to let you store all of your sampled sounds in a central integrated library where they reside together, regardless of source. An Alchemy sound file remembers its key range and what sampler it came from, and allows you to check or change its native sampler at any time. When you are editing a sound in the Alchemy environment, you can generally forget about its origin. All waveforms are edited using the same tools and processes.

Opening a Standard Sound File

- Choose the Open... command on the File menu.

A standard Macintosh Open dialog box appears.



- Click on the file you wish to open.
- Click on the Open button.

This will automatically put the sound in a waveform window on Alchemy's desktop, where you can edit as desired. Try this with one of the demo sounds supplied on the Alchemy sound disk.

When you are using the Open dialog box, Alchemy recognizes and will open a number of different file formats, including Sound Designer™, Dyaxis™, and the Audio IFF format (which is Alchemy's mono and stereo format of choice). For this reason, there is no need for an Import command. To make sure that a sound file from another source can be edited in Alchemy, just save it in any of the named formats.



Important Note: To open an SND Resource file (used by applications such as HyperCard®), choose the Import Resource command on the File menu. Then use the dialog box to select the application or document that contains the SND Resource you wish to open. A pop-up menu shows you the SND Resources that reside in any selected application or document. Use the pop-up menu to select the Resource you wish to open, then click on the Open button to open the file. It will appear in a new soundfile window.

Opening a Dyaxis Sound File



Opening Multiple Sound Files

- Choose the **Open...** command on the **File** menu.

A standard Macintosh Open dialog box appears.

- Click on the file you wish to open.

This will generally require that you use the dialog box to move to the Dyaxis' internal SCSI drive.

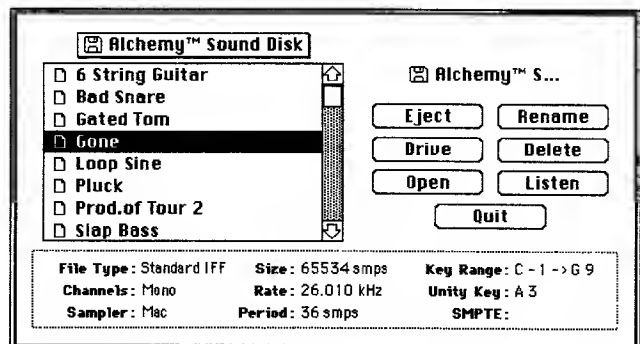
- Click on the **Open** button.

This brings up a dialog box warning you that Dyaxis files (which are usually very large) are edited directly on disk. That means that you should only edit copies of your original Dyaxis sound files, because disk-based operation actually changes the file on disk, not just in memory. If you edit your original Dyaxis files, you run the risk of damaging them with irreversible edits.

Important Note: The size of a selectable waveform range in a Dyaxis file depends entirely on the amount of memory available on your Mac. Alchemy will adjust any selected waveform range which is over the maximum to equal the maximum editable range.

- Choose the **Open Special...** command on the **File** menu.

Alchemy's Open Special dialog box appears.



Playing a Sound File Directly from Disk



Saving a New, Untitled Sound File

- Click on the file you wish to open.
- Click on the **Open** button.

This will automatically put the selected sound in a wave-form window on Alchemy's desktop. The Open Special dialog will remain on your screen to open more files.

- When finished, click on the **Quit** button.
- Choose the **Open Special...** command on the **File** menu.

Alchemy's Open Special dialog box appears.

- Click on the file you wish to open.
- Click on the **Listen** button.

The selected file will be played back directly from disk. Even files that will not fit in memory can be played back in 8-bit form directly from the hard disk. If you have a Sound Accelerator™ card, the sound will be played back in 16-bit stereo.

- When finished, click on the **Quit** button.

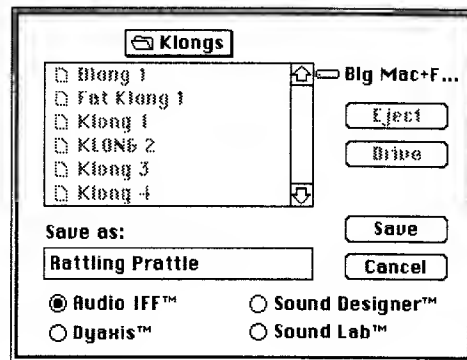
Note: Large sound files may only be played back directly from hard disks. Floppy disks are too slow to allow Listen button playback, except on short sounds.

- Choose the **Save As...** command on the **File** menu.

The Save As dialog box appears on your screen.

- Select the folder in which you want the file saved.
- Type a file name for the new file.

Updating an Existing File with Current Edits



- Choose the format in which you wish to save the file.

The Save As... command allows you to save the active file in any of five available formats (Sound Designer™, Dyaxis™, Sound Lab™, SND Resource, and Audio IFF). Audio IFF is the standard Apple interchange format, and the only one that saves your stored view memories, so it is recommended.

- Click on the Save button.

- Choose the Save command on the File menu.

The Save command automatically saves the currently selected file over its original source, and opens no dialog box.

Saving a Sound File Under a Different Name or Format



- Choose the **Save As...** command on the **File** menu.
- Select the folder in which you want the new version saved.
- Type the new file name.
- Choose the format in which you wish to save the file.

The **Save As...** command allows you to save the active file in any of five available formats (Sound Designer™, Dyaxis™, Sound Lab™, and Audio IFF). Audio IFF is the only one that saves your stored view memories, so it is recommended.

- Click on the **Save** button.

Important Note: To save the current soundfile as an **SND Resource** (for use in applications such as HyperCard®), choose the **Export Resource** command on the **File** menu. Then use the dialog box to select the application or document that should contain the **SND Resource**. A pop-up menu shows you the **SND Resources** that are already resident in the selected application or document. Click on the **Save Snd** button to save the file.


DAN: The Distributed Audio Network

Alchemy was designed as a central network controller which can get, edit, and send sampled sounds with many different samplers as the sources and destinations. To set up your network you'll need to tell Alchemy what samplers you have, and where they are in your communications chain. Once you've done this Alchemy will always remember your network when you start it up. As an Alchemy owner, you'll want to personalize Alchemy for your network as the first step in enabling Alchemy's capabilities.

Adding a Sampler to the Network

- Choose the New... command on the Network menu's Instrument pop-up menu.

The New Instrument dialog box appears. You'll be completing this dialog box one time for every sampler on your network



- Use the pop-up menu "Sampler ID" to select your sampler type.
- Click on the button in front of MIDI, RS-422, or SCSI to select your sampler's communication type.

For most samplers you will select MIDI.

- Click on the button in front of your sampler's communication clock rate.

For most samplers you will select 1 Mhz.

- Use the pop-up menu next to MIDI Channel to set your sampler's MIDI channel number.

If you're using a MIDI patcher:

- Click in the MIDI Patcher box and use the pop-up menu to set the MIDI channel that controls your patcher (Patcher Channel).

Deleting a Sampler from the Network

Editing a Sampler's Communication Settings

- Use the pop-up menu to set the **Patcher Program** number that will connect to the sampler you're adding.
- Click on the **New** button.

You have now successfully added a sampler to your network. You must complete this dialog box one time for each one of your network samplers. When you have done this, you will find a list of all your defined samplers at the bottom of the Network menu. If you have more than one of the same sampler type, Alchemy will automatically number them to help you keep track.

- Choose the name of the sampler you wish to delete on the bottom of the Network menu.

A check mark appears in front of the selected sampler.

- Choose the **Delete** command on the Network menu's **Instrument pop-up** menu.

You've just deleted an instrument from your sampler network. You'll notice that its name no longer appears at the bottom of the Network menu.

- Choose the name of the sampler for which you wish to change settings on the bottom of the Network menu.

A check mark appears in front of the selected sampler.

- Choose the **Edit...** command on the Network menu's **Instrument pop-up** menu.

The Edit Instrument dialog box appears. This box is nearly identical to the New Instrument box, except that it can only be used to change the settings for a sampler which is already on your network.

Getting a Sound from the Network

- Use the boxes, buttons and pop-up menus to change the sampler's network settings to reflect its new communications information.

See Adding a Sampler and the Reference section of this manual for more information about this.

- Click on the OK button.

The selected sampler is now defined with the new network information.

- Choose the New command on the File menu.
- Select the source node (sampler) from which you want to retrieve the sound.

There are two ways to do this. One way is to choose the sampler's name at the bottom of the Network menu. When you do so, a check mark appears in front of the sampler you've just chosen.

The other way to select a source sampler is in an actual waveform window. If you have set the Preferences... on the Action menu to show Instrument ID, the currently selected network sampler for any waveform window is displayed in its lower left corner. When a waveform window is selected, the sampler shown in its lower left corner becomes the active node. To select a different node (sampler), click the mouse inside the tiny window which shows the instrument ID. When the mouse cursor changes to an up or down arrow, use it to select another network node.

- Choose the Get Sound command.

This will open the keyboard dialog box which lets you select the instrument, preset or voice number, layer, and key range of the sound you want.

- Use the pop-up menu to select the instrument (preset) you wish to retrieve from.

Getting All Sounds from a Sampler

This is only applicable to certain samplers. See the sampler-specific information in the Appendix of this manual to learn more.

- **Use the pop-up menu next to “Layer:” to select the sound layer.**

This is not active for samplers like the Ensoniq Mirage or E-mu SP-1200.

- **Click the mouse on the keyboard to highlight the wave (key range) you wish to retrieve.**

You can also select the wave by using the pop-up menu next to “Sampler:” in the dialog box. Notice that the mouse remote-plays all keyboard layers under the highlighted key range on the selected sampler.

- **Click on the OK button.**

A thermometer-type display appears to let you monitor the progress of your transfer. After a short period the sound will appear on the screen in a waveform window ready for editing. Retrieving a stereo sound automatically requires one transmission pass for each channel, and the retrieved sound will appear in a stereo waveform window. The number of sounds which may be displayed and edited on the screen at any time is limited only by available memory.

- **Choose the New command on the File menu.**
- **Select the source node (sampler) from which you want to retrieve all sounds.**

There are two ways to do this. One way is to choose the sampler’s name at the bottom of the Network menu. When you do so, a check mark appears in front of the sampler you’ve just chosen.

Getting a Waveform Range from the Network



The other way to select a source sampler is in an actual waveform window. If you have set the Preferences... on the Action menu to show Instrument ID, the currently selected network sampler for any waveform window is displayed in its lower left corner. When a waveform window is selected, the sampler shown in its lower left corner becomes the active node. To select a different node (sampler), click the mouse inside the tiny window which shows the instrument ID. When the mouse cursor changes to an up or down arrow, use it to select another network node.

- **Choose the Get All command.**

This will open the keyboard dialog box which lets you select the instrument/preset from which all sounds will be retrieved. The Get All command will attempt to retrieve all sounds from the indicated instrument/preset. Each will appear in its own window. When memory is full, a dialog box will appear and let you know.

Note: Use the keyboard dialog's octave adjustment scale to slide the keyboard up or down so that it is a window into the correct sampler keyboard range. The octave scale lets you look at key ranges which are above or below a samplers normal play range, but which may still be played over MIDI.

- **Select a range in the active waveform window and choose the Get Range command on the Network menu.**

The Get Range command only works with sounds you've just retrieved from an Ensoniq EPS or Mirage, Roland S-50/S-550, or E-mu EIII. Use it to "re-get" a waveform range for editing purposes (if you accidentally erase something, for example). The Get Range command opens no dialog box, because the sound's source information has already been defined.

Sending Sounds to the Network

Replacing an Existing Sound


The possibilities for sending sounds to your network of samplers are much more varied than those for retrieving sounds. The following sections describe these different types of transfers and give detailed instructions for executing them.

- Activate the waveform window of the sound you wish to send.
- Choose the name of the destination sampler on the Network menu.

A check mark will appear in front of the sampler you select. Another way to select a destination sampler is in the actual waveform window. If you have set the Preferences... on the Action menu to show Instrument ID, the currently selected network sampler for any waveform window is displayed in its lower left corner. When a waveform window is selected, the sampler shown in its lower left corner becomes the active node. To select a different node (sampler), click the mouse inside the tiny window which shows the instrument ID. When the mouse cursor changes to an up or down arrow, use it to select a different destination node.

- Select the Send Sound command on the Network menu.

The Keyboard dialog box appears so you can choose the sound you will be replacing.



Assigning a New Voice and Key Range



- Use the pop-up menu to select the instrument (preset) you wish to assign the sound to.

This is only applicable to certain samplers. See the sampler-specific information in the Appendix of this manual to learn more.

- Use the pop-up menu next to “Layer:” to select the destination sound layer.

This is not active for the Ensoniq Mirage and E-mu SP-1200. You use it to assign more than one sound to be played back by the same key range. When you are sending stereo sounds, layer is disabled because Alchemy automatically uses both layers for the stereo image.

- Click the mouse cursor on the keyboard to select the playback key range of the sound you’re replacing.

Notice the the mouse remote-plays the selected sampler, so if you have sounds in your sampler memory, you’ll hear what you’d be replacing. You can also select the wave by using the pop-up menu next to “Wave:” in the dialog box.

- Click on the OK button.

A thermometer-type display appears to let you monitor the progress of your transfer.

Important Note: If you just want to return a sound to the same key range it originally came from, you don’t need to select any destination key range. All sounds retrieved through Alchemy and stored in the Audio IFF format will retain their original key range information. To find out more about this, see Building a Keyboard from Scratch, which follows.

- Activate the waveform window of the sound you wish to send.

- **Choose the name of the destination sampler on the Network menu.**

A check mark will appear in front of the sampler you select.

- **Select the Send Sound command on the Network menu.**

The Keyboard dialog box appears so you can set key range and other information for the sound you'll be sending.

- **Use the pop-up menus to select the instrument (preset) you wish to assign the sound to.**

This is only applicable to certain samplers. See the sampler-specific information in the Appendix of this manual to learn more.

- **Use the pop-up menus next to "Layer:" to select the destination sound layer.**

This is not active for the Ensoniq Mirage and E-mu SP-1200. You use it to assign more than one sound to be played back by the same key range. When you are sending stereo sounds, layer is disabled because Alchemy automatically uses both layers for the stereo image.

- **Click in the Assign New Voice box.**

The sound file's native key range is automatically highlighted and its native Unity key appears as a single dark gray key. If this is the range you want, you can click OK now.

- **Drag the mouse cursor across the keyboard to select the new key range.**

- **Click and drag the single dark gray key to the new Unity key position.**

The Unity key is the key which plays back the sample without transposing it up or down. This is where the waveform will sound exactly as it did when sampled.

Sending All Open Sound Files to a Sampler

A thermometer-type display appears to let you monitor the progress of your transfer. If assigning a new voice requires that you delete any sounds from your sampler's memory, a dialog box appears to ask your permission. In such a case, click OK to proceed.

- **Activate the waveform window of the first sound you wish to send.**
- **Choose the name of the destination sampler on the Network menu.**

A check mark will appear in front of the sampler you select. If you hold down the option key while selecting the destination sampler, all open windows will be set to that sampler.

- **Select the Send All command on the Network menu.**

The Keyboard dialog box appears so you can set key range and other information for the sounds you'll be sending.

- **Use the pop-up menus to select the instrument (preset) you wish to assign the sounds to.**

Only one keyboard map of sounds can be sent at once.

- **Click on the button in front of Retain, White, or Every at the bottom of the window.**

"Retain" will keep the samples' existing key range information. "White" will place all sounds on the white keys, starting from the bottom of the indicated keyboard (for drum sounds or sound effects). The "Every" setting allows you to assign each sound to a specific number of keys, starting at the bottom of the indicated keyboard. If you choose "Every," you must also type in a number of keys to be used for each sound.

Sending a Waveform Range

- Select a range in the active waveform window and choose the Send Range command on the Network menu.

The Send Range command only works with sounds you've just retrieved from an Ensoniq EPS or Mirage, Roland S-50/S-550, or E-mu EIII. Use it to return a waveform range to its original source (to test out a quick fix from a keyboard, for example). The Send Range command opens no dialog box, because the sound's source information has already been defined.

Sending Stereo Sounds

- Follow exactly the same procedure you would use to send a mono sound.

The only differences you'll notice when sending a stereo sound is that all layer information is disabled, and twice as much memory is required at the destination sampler. The layers are disabled in order to automatically assign both channels of the stereo image to the same key range, and the extra memory requirement is simply for the extra channel of wavedata.

When you've finished setting up your key assignment, and have clicked OK, two thermometer-type displays appear sequentially to let you monitor the progress of both channel transfers.

Time Domain Processing

Waveform Window Display Tools

reverse or invert sounds, band effects, mix sounds together, fade and crossfade sounds, etc. These are all processes which adjust a waveform's orientation, sample order, or amplitude, and as such they open up a world of sound design possibilities. This section of the manual will help you to develop an understanding of some of these waveform editing techniques. If any of the concepts confuse you, review the About Sound section of this manual, where the conceptual basis of time domain processing is discussed.

Once you understand how to use Alchemy to edit waveforms in the time domain, you can progress to the explanation of Alchemy's harmonic spectrum analysis and resynthesis. If you are curious about particular palette tools, you can find very specific palette information in the Reference section of this manual. As you go through this section, try out the different tools on sounds of your own. This should help you develop your own style of working with Alchemy.

The waveform window display is really Alchemy's center of activity. It can show you a picture of your entire sampled waveform, of a single waveform period, or anything in between. It's the environment you'll use to view and edit any sound over time. You can think of the waveform window as a fully adjustable sound microscope. If you learn to use it well, you'll be able to make Alchemy truly shine.

Any time you open a sound file, or get a sampled sound from a sampler, Alchemy opens a waveform window to show you that sound. When a waveform window is opened, it is automatically sized to show you the entire waveform. Because a sampled sound contains so much information, viewing an entire waveform may not be the best way to see what you're working on. By manipulating the waveform window display, you can look at any part of any waveform at many different resolutions. This allows you the flexibility to select everything from an entire waveform to a single sample for editing.

Standard Waveform Window Modes

Once you've used the waveform window display to view and select an area for editing, you can execute most process commands from the palette which is always visible on the left side of the screen. Tasks like reversing, fading, or inverting a sound only require one click. Even the most involved processes take no more than three steps. Everything happens right before your eyes.

Whenever you start up Alchemy and get or open a sound, the active waveform window will automatically be in Selection mode. This is Alchemy's default mode, and it allows you to select waveform ranges for playback and editing. You'll know you're in Selection mode, because the Selection icon (at the upper left corner of the palette) will be highlighted.

Alchemy also has four other modes, Waveform Draw mode, Loop Splice mode, Amplitude Envelope mode, and Frequency Envelope mode. However, these have very specialized uses which you'll read about later in this section.

The Selection Mode

Before any process may be executed, you must select the waveform range over which the process is to take place. For this reason it is important to familiarize yourself with the tools explained over the next few pages. They are the only ones you will use to select ranges for editing.

Selecting a Waveform Range


- **Open any sound file.**
- **Place the mouse cursor over the active waveform.**

The active window is indicated by a highlighted title bar. When the cursor is over the waveform in this window, the mouse arrow changes to an "I-beam." If it doesn't, you're not in Selection mode.

- **Click and drag the cursor to the left or right in the waveform.**

As you do this, the range you are selecting will be highlighted in black. If you look at the palette's numeric display (located at the bottom of the palette), you'll see changing numbers which show the position of the I-beam, the value of the sample at that position, and the size of the selected range. (For more information, see the Numeric Display in the Reference section of this manual).

- Click on the palette's Speaker icon or press the space bar to play back the selected range.



Once a range is selected, it can also be processed. That means that any selected range can be cut, copied, reversed, inverted, faded, etc.

Sliding a Selected Range

- Hold down the option key and place the mouse cursor over the selected waveform range.

The mouse cursor changes into a "grabber" hand.

Extending a Selected Range

- While holding down the option key, click and drag the whole selection range to a new location.

- Hold down the shift key and click the mouse cursor outside of the selected waveform range.

The current selection range is extended to the new point.

- While holding down the shift key, click and drag the mouse cursor to adjust the extension of the currently selected waveform range.

Creating a Stereo Soundfile

- Open a new or existing sound file using the New... or Open... commands on the File menu.

- Choose the Mono to Stereo command on the File menu.

The waveform window will reappear with an upper (left channel) and a lower (right channel) waveform. These waveforms start out identical, but they can be edited together or separately.

Selecting a Single Channel of a Stereo Display

- Place the I-beam cursor over either the upper (left channel) or lower (right channel) waveform.
- Click and drag to select a range, just as you would in a mono display.

If you want to place the blinking insertion point, just click the mouse, instead of clicking and dragging.

Selecting Both Channels of a Stereo Display



One of Alchemy's distinctions is that can be used to create stereo sounds from mono waveforms, and to edit existing stereo sounds. Stereo sound files may consist of two stereo channels containing a true stereo-sampled image (from a Dyaxis system, for example), or they may be created within Alchemy to build stereo image effects, or match up mono-sampled sides of a stereo image. Once a stereo sound file is created, it can be stored and retrieved as such in the Audio IFF format.

When you send a stereo sound file to a sampler, Alchemy automatically assigns both channels to the same key range, places them on different layers, and pans them left and right correctly. True stereo playback can be accomplished on all network samplers except the Ensoniq Mirage and E-mu SP-1200. To find out more, see the Applications at the end of this section and the sampler-specific information in the Appendix of this manual.

- Place the I-beam cursor over the channel separator (the line separating the upper display from the lower display).
- Click and drag to select a range, just as you would in a mono display.

If you want to place the blinking insertion point, just click the mouse over the channel separator, instead of clicking-and-dragging.

Important Note: You can view either the left or right channel alone by choosing either the Left Channel or Right Channel command on the Channel Display pop-up menu located on Windows menu. The channel you select will automatically be sized to fill the waveform window. To return the waveform window to its stereo display, choose the Stereo command on the Channel Display pop-up menu. A check mark always appears on this menu to tell you what you're looking at.



Important Note: Playback of stereo sounds depends on your computer and on your samplers. At this time, only Macs with a Sound Accelerator™ card can play back sounds in stereo. If both channels are selected, on any other Mac, only the left (upper) one plays.

Selecting a Loop Range

- Choose the Select Loop command on the Edit menu.

The current sound file's loop range is automatically selected for editing.

Placing an Insertion Point

- Point the mouse I-beam anywhere in the current waveform and click once.

When the insertion point is flashing in a waveform window, the Speaker icon will play back the entire sampled sound and its loop. The insertion point may be used instead of a selection range to mark a Paste, Mix, or Insert range start. When you are pasting, mixing, or inserting sounds, you may want to place the insertion point at the point where the paste or mix should begin. This assures that all of the Clipboard contents will be pasted or mixed. (Pasting or mixing with a range selected just fills up the range.)

Saving Selection Ranges

- Adjust the current waveform window so that it is displaying the selection range (and view) you want to save.



- Click on the palette's SAVE button.
- Click on the view number you want for the range and view you're saving.

The next time you click on that number, the view and range you just saved will automatically be displayed. The saved view will remain under that number until you replace it with another view.

You will find the view memories to be extremely useful in sound design. You can use them to keep track of fade or crossfade areas, store potential loop ranges, remember mixing ranges, etc. By using the view memories carefully, you'll be able to avoid the repetitive task of readjusting the waveform window's resolution every time you want to look at an often-used waveform area.



Important Note: If you want to make sure your waveform view memories are saved with their sound file, you must save your sound files in the Audio IFF format. This is the only format which retains view memory information.

Recalling Saved Selections

- Click on the waveform view number in the palette's view memories box.

The waveform view and selected range appear immediately in the active window. Each sound file allows eight separate view memories.

Sound File Navigation

The Overview Display

Creating a Sound File Overview


Alchemy is equipped with a number of navigation tools which make it easy to adjust the current waveform view to many resolutions with few steps. These tools may be divided into three types: The Overview display, the Locator icons, and the Zoom functions. For explanations of each, just keep reading.

The Overview display is one of the most useful of the waveform window display tools. The Overview display is a navigation tool which allows you to move around very quickly in any waveform. Picture it as a map of the entire sound you're working in; you can use it to move from one waveform area and resolution to another in one step.

- Click on the Overview icon.



You can accomplish the same thing by choosing the Show Overview command on the Windows menu. The overview display splits the current waveform window to show your waveform from two different views. The larger, lower wave display is the working area where you'll be doing your editing. It functions exactly like the normal waveform display, and can be used for all editing and looping tasks.



whole waveform. You can also use it as a quick and accurate navigation tool to move all over the waveform you're working on.

Navigating from the Overview Display

- **Drag a dotted sizing box with the mouse cursor over any waveform range in the upper (overview) display.**

The waveform area you just selected in the overview display will automatically be sized to fit in the lower waveform display. The waveform section which currently appears in the lower waveform display is indicated in the overview display by an empty rectangular bar above the overview waveform. The currently selected waveform range is indicated in the overview display by a greyed area, or 'gel.' It's important that you remember the difference between these two indicators. The short rectangle indicator shows you what you're looking at in the lower display, while the gray area shows you what waveform range is actually selected for editing.

Locating in the Overview Display

- **Once a waveform range has been selected in the overview display, click the mouse once within the overview display to view any other waveform section at the same resolution.**

When you do this you are simply moving your window on the waveform to another range. This does not affect what is currently select, it only adjusts the view.

Changing Overview Resolution

- **Click on the Snapshot icon.**





The Snapshot icon (or Take Snapshot command) updates the overview display with a copy of the current waveform view. The new overview display may be at a different resolution, but you can still use it as before for navigation.

If the active waveform window is not showing an overview display, this command opens one and places in it a copy of the current waveform view.

Important Note: You can reset the overview display to show the entire waveform again by clicking on the Overview icon a second time.

Locator Icons

Below the waveform view icons you see the four Cursor Locator icons, which are grouped into a pair of I-beams and a pair of loop cursors. These icons allow you to move immediately to the beginning or end of the current selected range or loop, without changing your magnification level.



The Locator Icons

Centering Your View on Range Start

- Click on the left Cursor Locator icon.

The left I-beam icon centers the display on the insertion point (or on the range start point, if a range is selected). However, it doesn't change the view magnification.

Shortcut: Holding down the command key when you click any of the Cursor Locator icons automatically zooms full in to the selected point.

Centering Your View on Range End

- **Click on the right Cursor Locator icon.**

Clicking on the right I-beam icon automatically centers the view in your waveform window on the end of the currently selected range, but it doesn't change the view magnification.

Centering Your View on Loop Start

- **Click on the left loop cursor locator icon.**

Clicking on the left loop cursor icon automatically centers the view in your waveform window on the current sound file's loop start point, but it doesn't change the view magnification.

Centering Your View on Loop End

- **Click on the right loop cursor locator icon.**

Clicking on the right loop cursor icon automatically centers the view in your waveform window on the current sound file's loop end point, but it doesn't change the view magnification.

Changing Waveform Resolution (The Zoom Functions)

Aside from the overview display, which may be the single most powerful zoom tool, there are a host of other zoom functions. There are two Zoom icons located at the center of the palette. The one on the left is the Zoom In icon, and it looks like a magnifying glass focusing on a large triangle wave. The Zoom In icon is the tool you'll use to move the current view inward, which fits less of the waveform in the window, but shows it to you at a greater resolution. When the waveform view is zoomed all the way in, the display resolution is at its highest, and each sample is displayed as a single pixel.

Zooming In on a Waveform

The Zoom Out icon is located to the right of the Zoom In icon, and you can use it to zoom the view of the current waveform outwards, thereby fitting more of the waveform into the window. When the waveform view is zoomed all the way out, the entire waveform in memory will be visible.

To zoom in one step at a time:

- Click once on the Zoom In icon.



This zooms in at the left edge of the current waveform window by a factor of two until each pixel represents a separate sample. That is maximum zoom. You can also accomplish this by choosing the Zoom In command on the Action menu.

To zoom all the way in:

- Hold down the command key and click on the Zoom In icon.

This zooms in all the way at the left edge of the current waveform window. You can also accomplish this by choosing the Zoom Full In command on the Action menu.

Zooming Out on a Waveform

To zoom out one step at a time:

- Click once on the Zoom Out icon.



This zooms out at the left edge of the current waveform window until the entire waveform is visible in the waveform window. You can also accomplish this by choosing the Zoom Out command on the Action menu.

Automatically Zooming into a Selected Range

Customizing Your Environment

Setting Basic Preferences

To zoom all the way out:

- **Hold down the command key and click on the Zoom Out icon.**

This zooms out all the way to show the entire current waveform. You can also accomplish this by choosing the Zoom Full Out command on the Action menu.

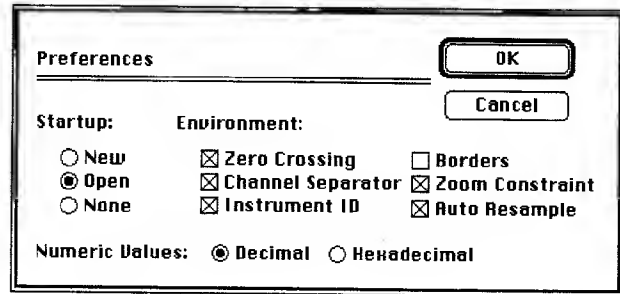
- **Click on the palette's Fit Selection icon.** 

The fit selection icon automatically sizes the currently selected waveform range to fit exactly in the waveform window. This is the most efficient way to zoom in to whatever is selected, and is an easy way to get the closest look possible at the range where edits can be executed. You can also accomplish this by choosing the Fit Selection command on the Action menu.

Alchemy contains a number of functions which allow you to personalize your system so that it functions the way you like. Here you'll find a number of variables you can set to make Alchemy look and operate the way you prefer.

- **Choose the Preferences... command on the Action menu.**

This opens the preferences dialog box, which lets you set Zero Crossings, Channel Separator, Overview Separator, White Space, Zoom/Tile Constraint, and Sampler ID. Here's what each do:



Zero Crossings: Displays a black line representing the zero line in waveform windows.

Channel Separator: Displays a black separator line between right and left channels in stereo waveform displays.

Borders: Leaves open space above and below waveforms to show tabs at loop cursor bottoms and to facilitate easy grabbing of threshold bars.

Instrument ID: Shows an adjustable display in the lower left corner of all waveform windows which lists their current network sampler source/destination.

Zoom Constraint: Prevents the Tile and Strip commands from placing waveform windows underneath the palette.

Auto Resample: Automatically prompts you when you send a sound to choose a sample rate that is compatible with the destination sampler. The sound is then resampled to the compatible rate before it is sent.

The preferences dialog also lets you to choose decimal or hexadecimal values for sample display, and allows you to elect what window will be opened upon program startup.

- Click on the buttons in front of all desired preferences.
- Click on the OK button.



Choosing Axis Units

Viewing a Waveform with Axis Markers



Setting the Macintosh Sample Playback Rate

All preferences will take effect immediately, and will be remembered when you quit the program.

Note: The Preferences dialog can also be displayed by holding down the Command key, and clicking on the Speaker icon.

- Use the **Axis Units pop-up menu on the Action menu.**

Choose either Samples, Seconds, or SMPTE. *Note:* The SMPTE format is set to standard 30 frames. From now on, whenever you click the axis markers icon, or choose the Show Rulers command, all axis markers will be displayed in the units you've chosen.

- Click the axis markers icon located near the top of the palette.

You can accomplish the same thing by choosing the Show Rulers command on the Action menu.

You can view the waveform window display with or without axis markers. The horizontal axis markers are displayed by sample number, seconds, or 30-frame SMPTE. The vertical axis is displayed in percent of maximum amplitude.

- Use the **Audio Output pop-up menu on the Action menu.**

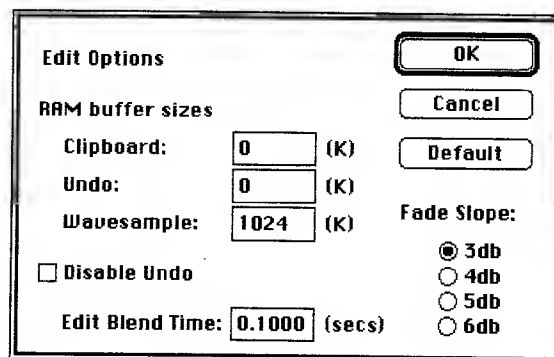
Choosing Mac fixed rate plays everything back at the Mac's set rate of 22.254 kHz, which assures you of the best fidelity, but doesn't play back at the right pitch. Choosing Mac Variable rate gives you a true picture of all sampled pitches, but loses fidelity by trying to play back rates out of

Editing Memory Use and Splice Options

its range. Most people prefer the wavesample's playback rate, because it seems truer to the original sound. The Sound Accelerator option is only available to those with Digidesign's Sound Accelerator™ card installed.

- Choose the **Edit Options...** command on the **Edit** menu.

This opens up the edit options dialog box, where you can allocate memory to certain editing functions, and adjust the automatic crossfade size for the Blend function.



- Type in your preferred memory amounts.

Hard disk owners should set the Clipboard and Undo sizes to 0, and leave the wavesample setting at the default. This will free up as much memory as possible for wavedata. Set the Blend Amount to indicate the size of the crossfade overlaps that will be used when editing is done with the Blend function turned on. See *Blending on Edits*, later in this chapter.

- Click on the **OK** button.

Waveform Editing Functions

Cutting a Waveform

Copying a Waveform

In Alchemy you accomplish basic editing using the same editing tools you'll find in almost every Macintosh program. That means that if you're familiar with cutting, copying, and pasting, there will be no surprises. Still, since Alchemy works with waveforms and frequencies, the way to use these standard Mac commands may need some introduction. Here is some information about the workings of Alchemy's basic editing commands. As usual, it will probably help you to try out these commands as you read about them.

- **Select the range to be cut in the active waveform window.**
- **Choose the Cut command on the Edit menu.**

The Cut command functions as you would expect. When you choose the Cut command from the Edit menu, the currently selected waveform range is cut from the waveform and placed on the Macintosh Clipboard for later pasting, mixing, inserting, or replicating. The entire waveform to the right of the cut range slides to the left to fill the opened space. The cut waveform remains on the Clipboard until you cut or copy something else, or shut down your Mac.

- **Select the range to be copied in the active waveform window.**
- **Choose the Copy command on the Edit menu.**

When you choose Copy from the Edit menu, an exact copy of the currently selected waveform range is placed on the Clipboard for later pasting, mixing, inserting, or replicating. No change takes place in the active waveform. As with the Cut command, a copied waveform remains on the Clipboard until you cut or copy something else, or shut down your Mac.

Pasting a Waveform

To constrain the paste to a selected range:

- **Select a destination waveform range.**
- **Choose the Paste command on the Edit menu.**

The Paste command takes the waveform range which is currently on the Clipboard and pastes it into the active waveform. If a waveform range is selected, the Clipboard waveform is pasted only over the selected range. If the Clipboard waveform is longer than the selected paste range, only the beginning of the Clipboard waveform is pasted. It's important to remember that a pasted waveform always covers up the waveform underneath it. If you want to add the Clipboard waveform to the one in the waveform window, use the Mix command, explained below.

To paste the entire Clipboard from an insertion point:

- **Click the mouse to place the blinking insertion point at the paste start point.**
- **Choose the Paste command on the Edit menu.**

When the active waveform contains the blinking insertion point, the entire Clipboard waveform is pasted over the existing waveform, beginning at the insertion point position.

Mixing a Waveform

To constrain the mix to a selected range:

- **Select a destination waveform range.**
- **Choose the Mix command on the Edit menu.**

The Mix command operates much like the Paste command, but instead of pasting the Clipboard waveform over a waveform range in the waveform window, it mixes the Clipboard waveform and the waveform window range together, producing a combination of the two sounds. The mix covers only the selected destination range, regardless of the size of the Clipboard.



"Out-point" Editing

To mix the entire Clipboard from an insertion point:

- **Select an insertion point where the mix should begin.**
- **Choose the Mix command on the Edit menu.**

When the insertion point is active, the Mix command adds the entire Clipboard to the current waveform, starting at the insertion point.

Important Note: There is one thing you should consider when mixing sounds. Mixing actually adds two sounds together, so if they both are waveforms with relatively high amplitudes, their added amplitudes may exceed the maximum at some points. This results in a clipped signal, which you'll usually want to avoid. To make sure that you don't exceed the maximum amplitude when mixing waveforms, you may want to scale the amplitudes of the waveforms before mixing them. This is called pre-scaling. For more information about this, see Scaling and Volume Smoothing, later in this chapter.

To Paste or Mix the Clipboard starting at range end:

- **Select a destination waveform range.**
- **Hold down the option key while you choose the Paste or Mix command on the Edit menu.**

This version of the Paste and Mix commands takes the waveform range which is currently on the Clipboard and pastes or mixes it into the active waveform beginning with range end and Clipboard end. This is the only way to match up the end or "out" points of Pasted or Mixed waveforms.

Inserting a Waveform

- Place the blinking insertion point at the desired insert point.
- Choose the **Insert** command on the **Edit** menu.

Alchemy's **Insert** command takes the waveform which has been cut or copied to the Clipboard and inserts it after the insertion point in the active waveform window. When you use the **Insert** command, you'll usually be making the overall waveform longer. In such a case, a dialog box always appears to ask you if it's OK to increase the overall length of the sound by inserting the new wave section. If you choose **Yes**, the total waveform length will be increased so that the Clipboard waveform can be inserted. If you choose **No**, the Clipboard waveform will be inserted anyway, but the total length of the waveform will not be increased. This is accomplished by cutting off the tail end of (truncating) the active waveform to squeeze the Clipboard waveform in. If you choose **Cancel**, the insertion will not be made.

Blending on Edits


- Set the desired **Blend Amount** using the **Edit Options...** command on the **Edit** menu.
- Choose the **Blending** command on the **Edit** menu.

The **Blending** command is actually a function that can be toggled on and off. When **Blending** is turned on, the beginning and ending points of all **Cuts**, **Pastes**, and **Inserts** will automatically be crossfaded with the existing waveform. Here is an illustration of the effect **blending** has upon a standard cut:

Clearing a Waveform

- Select the range to be cleared.
- Choose the Clear command on the Edit menu.

The Clear command clears all samples out of the selected waveform range, which erases any sound in that range. The Clear command works by setting all samples in the selected



The size and slope of this auto-crossfade (Blend Amount) is set using the Edit Options... command. Among other things, this function guarantees that no clicks or pops will be created during an edit by making sure that no instantaneous amplitude changes are caused at edit splice points. This is a very powerful tool, with far-reaching sound design implications. Although small blend ranges are generally used most often, try pasting a Clipboard waveform into a new sound with a very long Blend Amount setting. Remember, blending often will not maintain the overall duration of a sound, because the blend range crossfades overlap.

Extracting a Waveform

Undoing a Process

range to zero. It does not cut the selected range to the Clipboard, and does not change the sound file's length in any way. If you want to shorten the length of a sound file by actually cutting off its end, you'll need to use the Extract command.

- **Select the waveform range you wish to keep.**
- **Choose the Extract command on the Edit menu.**

Extracting is a clever way of truncating. By selecting the waveform range you wish to keep, you can truncate either the beginning or end of the current sound file. In fact, by using the Extract command you can truncate both the beginning and end at the same time. When you extract a waveform range, Alchemy redefines the present sound file to contain only that range and gets rid of everything else. The next time you save the extracted file, it will replace the old version. Remember that you can use the Undo command to reverse the process or the Revert command to return to the last saved version if you execute an extract by mistake.

- **Choose the Undo command on the Edit menu.**

The Undo command reverses the last command or process you executed. It can only remember one step back, so be careful if you are depending on the Undo command to cancel a finished operation. If you accidentally choose another command or process before undoing, you will be unable to backtrack far enough to reverse previous steps.

If you use it carefully, the Undo command can be excellent for "before and after" comparisons, but remember, you should always save the waveform you're working on so you can revert back to it. Otherwise you run the risk of damaging the waveform by executing two processes, and only being able to undo one of them.

Time Domain Processing Functions

Along with the standard Macintosh editing tools, Alchemy offers a number of advanced editing functions which are uniquely suited to the manipulation of sound. All of these editing tools can be chosen from Alchemy's palette, and most are also on the Process menu. Most of these advanced editing functions have been designed to allow one-step execution of often-used processes like waveform reversal or replication. Learning to use these tools will greatly decrease the time you'll need to complete many sound editing tasks.

Fades and Crossfades

Alchemy allows you to accomplish most standard fades and crossfades from the palette with a single mouse click, and the actions take place directly in front of you. Aside from offering immediate visual feedback, a variety of fade slopes are also available, and can be set at any time.

Setting the Fade Slope

- **Choose the Edit Options... command on the Edit menu.**

This opens the Edit Options dialog box. Use the radio buttons at the bottom of the box to choose a 3dB, 4dB, 5dB, or 6dB fade slope.

- **Click on the button in front of the fade slope you desire.**

For some general information about fade slopes, see the Looping Workshop which is Application 1 at the end of this section of this manual.

- **Click on the OK button.**

The fade slope you just chose is now Alchemy's default fade slope. It will automatically be used for all fades and crossfades until you change it.

Fading In a Waveform

- Select the range to be faded in.

- Click on the Fade In icon on the palette.



This same function can be executed by choosing the Fade In command on the Process menu. As you can see, fading a waveform range in or out is a simple process. The fade slope for both fade in and fade out defaults to whatever curve you've chosen using the Fade Options... command on the Process menu. The fade options can be set by holding down the command key and clicking on the Fade In (or any Fade or Crossfade) icon.

Fading out a Waveform

- Select the range to be faded out.

- Click on the Fade Out icon on the palette.



This same function can be executed by choosing the Fade Out command on the Process menu. Fading out a waveform range functions almost exactly like fading one in. The fade slope once again defaults to whatever curve you've chosen using the Fade Options... command on the Process menu. The fade options can also be set by holding down the command key and clicking on the Fade out (or any Fade or Crossfade) icon.

Crossfading Two Waveforms

- Select the waveform range to be faded out.

This will be the beginning (attack) portion of the final sound.

- Choose the Copy command on the Edit menu.

The selected waveform range is now on the Mac Clipboard ready for crossfading.

- **Select the target waveform range which will be faded in.**

This will be the ending portion of the final sound.

- **Click on the Crossfade icon.**



Crossfading is the process of mixing one sound as it fades out, with another sound as it fades in. The result is a smooth transition from one sound to another (the concept of crossfading is explained in the About Sound section of this manual). The fade slope used for crossfading is the same one you set using the Fade Options... command on the Process menu. It can also be set by holding down the command key and clicking on the crossfade (or any fade) icon.

Note: If you hold down the option key when you click on the Crossfade icon, the Clipboard waveform will be faded in at the end of the selected range, rather than faded out at the beginning of the selected range.




Amplitude Scaling

Amplitude scaling a sound is the process of increasing or decreasing its amplitude evenly over its duration. It is a particularly useful tool when you are mixing sounds together, because it allows you to adjust exactly their relative levels.

You probably remember the discussion of clipping under the Mix command earlier in this section. Mixing two sounds which both have relatively high amplitudes often produces amplitudes at certain points which are above the maximum allowable amplitude. This causes all samples which are above maximum to be set to the maximum, which essentially clips the waveform at that point. To avoid this problem simply scale down the amplitude of one or both sounds so that the product of their mixture will never exceed the maximum allowable amplitude.

Scaling a Waveform



- Select the Auto Zero command on the Edit menu.

It's always a good idea to make sure that you select any waveform range to be scaled with the Auto Zero command checked on the Edit menu. This will insure that your scale won't cause an inconsistency in the middle of your waveform which will result in a click or pop.

- Select the waveform range to be scaled.
- Click on the Threshold Bars icon on the palette.

This displays the dotted threshold bars at the top (and bottom) of the active waveform window. You will use these bars to set the scaling threshold.

- Place the mouse cursor over the upper threshold bar until it turns into an up and down arrow.
- Click and drag the threshold bar to the desired threshold amplitude.

This threshold will be the amplitude maximum for the selected waveform range. If the threshold is below the highest amplitude in the waveform range, you'll be scaling down (decreasing amplitude). If it is above the highest amplitude you'll be scaling up (increasing amplitude). To check the actual factor by which you'll be scaling, look in the palette's numeric display as you drag the threshold bars. The number next to "Scale:" tells you what's happening. 1.0 is no change, 0.50 would be a scale to 50% of the former amplitude, and 2.00 would be a doubling of the former amplitude.

Note: If you have trouble grabbing the threshold bars, try turning on the Borders preference using the Preferences... command on the Action menu. This will give you more grabbing room.



Scaling Example

- Click on the Scale icon to execute the scale.

Your amplitude change has now been made.

Here's a short demonstration to help you understand how you might use scaling to create particular sound mixtures. Let's assume that you want to create a sound which is 60 percent human voice and 40 percent cement mixer. You might follow this procedure to do it:

- Open both a human voice sound file and a cement mixer sound file.
- Click on the human voice window to make it the active waveform.
- Select the entire waveform by double clicking on the Selection mode icon.
- Click on the threshold icon to show the horizontal dotted threshold bars in the waveform window.
- To set the threshold bars, place the I-beam cursor over either the top one or bottom one.
- Now click and hold the mouse button.

You'll notice that the palette's numeric display changes to show new information. Beg: and End: indicate the first and last sample of the selected scale range, and Scale: tells you the amount of waveform scaling the present threshold position will cause. 1.0 means no scaling, 0.75 means scale down to 75% of original amplitude, 1.33 means scale up to 133% of original amplitude, etc.

- Since you know you want your final sound to be 60% human voice, drag the threshold bar down until Scale: shows 0.60 in the numeric display.



Inverting a Waveform



- To execute the amplitude change click on the Scale icon.

The Scale icon is shaped like a balance and located toward the middle of the palette. Your waveform will now be at 60% of its former amplitude.

- Click on the cement mixer waveform window to make it active.
- Follow the same scaling procedure you just went through, but scale it down to 40% of its original amplitude (Scale: 0.40).
- Select and copy the entire cement mixer waveform to the Clipboard using the Copy command.
- To complete your mix, click on the human voice waveform window to make it active again.
- Make sure the entire human voice waveform is selected, and choose the Mix command.

This procedure would create a mixture of the 60% human voice and the 40% cement mixer.

Note: A sound's frequency content often contributes as much to its perceived "presence" as its amplitude. For this reason, some sounds with lower average amplitudes may seem to be louder than sounds with higher average amplitudes.

- Select the range you wish to invert.
- Click on the palette's Invert icon.

Inverting any waveform range turns it upside down, which is an important function in a number of looping methods (mirror looping in particular). Waveform inversion is accomplished by making all positive sample values negative, and all negative ones positive. Although inversion does change the orientation of the waveform range, it will not alter the waveform's sound in any perceptible way.

Reversing a Waveform



- Select the waveform range you wish to reverse.
- Click the palette's **Reverse** icon.

When you reverse a waveform you are reversing the order of the samples in the selected range. When you click on the Speaker icon (or press the space bar) the sound will be backwards. Repeat the process to return it to normal.

Replicating a Waveform



- Select the waveform range you wish to replicate.
- Choose the **Copy** command on the Edit menu.

A copy of the waveform range you wish to replicate is now on the Clipboard.

- Select the range over which you want to replicate the waveform you just copied.

This should preferably be at least a few times longer than the waveform you wish to replicate.

- Click on the **Replicate** icon

The previously selected range is now filled with copy after copy of the waveform range you wanted to replicate. This can be a particularly useful function for creating long sounds from single wave cycles and digital effects like echo.

Waveform Draw Mode

Drawing a Waveform




Waveform draw mode is one of the three modes available for any waveform window. It allows you to use the mouse cursor as a pencil to draw new wave shapes and redraw existing ones. You can use Waveform Draw mode to hand draw waveforms at any resolution. Wavedata is created or changed immediately and can be heard by switching back into Selection mode,

- Open any sound file.
- Adjust the waveform view to the resolution you prefer for drawing.

It will probably be best if you're zoomed all the way in.

- Click on the Waveform Draw mode icon (pencil) at the top of the palette.

When you do this the waveform window display will not seem to change, but the mouse cursor will be a pencil instead of a selection I-beam when it is over any window range.



Amplitude and Frequency Enveloping Modes



- **Place the mouse cursor over any range within the active waveform window.**

It turns into a pencil for drawing.

- **Click and drag to draw**

It is usually best to draw only at a very high resolutions (zoomed all the way in), because this allows you much greater accuracy.

- **Click on the Selection mode icon to return to Selection mode.**

This will allow you to work with what you've drawn. One interesting experiment is to zoom all the way in, draw a waveform, then loop it and listen to what it sounds like. Generally speaking, however, Waveform Draw mode is best used as a corrective tool for removing glitches, smoothing loop points, etc.

Important Note: If you want to draw a waveform in a new file, you must make sure that the file is not empty. In other words, it must contain more than zero samples for which you'll be drawing values. To adjust the number of samples in a new (or any) file, use the Soundfile Setup... command on the File menu.

Amplitude and Frequency Enveloping modes are two very powerful synthesis environments only available in Alchemy. Amplitude Enveloping Mode allows you to adjust or trace the amplitude envelope of any sampled sound. Traced envelopes can be copied and superimposed over other sampled sounds, or even pasted as waveforms for storage in an envelope library. Frequency Enveloping mode allows makes it possible to modulate the frequency of any sampled sound by drawing in a modulation envelope, or by pasting in any waveform to act as a modulation envelope. Alchemy does not distinguish between envelopes and normal waveform information, so either may be used as the modulator. To see what new synthesis doors this opens, read through the following section.

Amplitude Enveloping Mode



- Open the sound file containing the waveform whose amplitude envelope you'd like to change.
- Click on the Amplitude Envelope Mode icon.


You are now in Amplitude Enveloping Mode. The waveform will appear gray, and a horizontal amplitude envelope line will appear near the top of the Y axis. The blank envelope line is now set to full on (attack).

- Click the mouse cursor anywhere on the anywhere on the amplitude envelope line to create a break point.

A break point looks like a small black square, called a "knob."

- Drag the break point knobs to positions that define the envelope you wish to create.

The amplitude profile you define will be the one your sound will eventually have.



Frequency Enveloping Mode

- Click on the **Amplitude Fit** or **Amplitude Scale** icons.

Amplitude Fit will shrink and/or grow overall amplitude values to fit exactly under the new envelope. Amplitude Scale can only decrease amplitude values, and will adjust your waveform to the correct shape, but not necessarily to the same actual amplitudes. (See the next page for more information).


- Open the sound file containing the waveform whose frequency envelope you'd like to modulate.

- Click on the **Frequency Envelope Mode** icon.



You are now in Frequency Enveloping Mode. The waveform will appear gray, and a horizontal amplitude envelope line will appear on the X axis. The blank envelope line is now set to no modulation.

Frequency Envelope (pasted from Clipboard)



Scale (in semitones)

Editing Envelopes with Pencil and with Knobs



- Use the Process menu's Frequency Range pop-up menu to set the scale for the modulation you will perform.

The selected range defined how far the pitch will be modulated at full positive and full negative envelope values.

- Click the mouse cursor anywhere on the frequency envelope line to create a break point.

A break point looks like a small black square, called a "knob."

- Drag the break point knobs to positions that define the envelope you wish to create.

The frequency profile you define shows the way in which the sound's pitch will be changed when you issue the Frequency Mod command.

- Click on the Frequency Mod icon.

The pitch of your waveform will be modulated (shifted) according to the modulation curve defined by the modulation envelope.

Both amplitude and frequency envelopes can be edited in two ways. In their default modes, envelopes are edited with break point "knobs," which are created and dragged by clicking the mouse on the envelope lines. It is also possible to hand-draw envelopes to your liking using a pencil tool. Here's how to do it:

- Click on the Draw/Knob icon to toggle between envelope draw and envelope knob modes.

When draw is selected, the cursor turns into a pencil while over the waveform window. Edit envelopes by dragging an envelope shape. Both pencil and knobs can be used to edit either type of envelope.

Tracing Envelopes



A Note on Copying and Pasting Envelopes

Alchemy offers a powerful alternative to the creation of amplitude and frequency envelopes from scratch. The amplitude envelope of any existing waveform can be auto-traced for editing or superimposition over other waveforms. To do this:

- **Open the sound file containing the waveform whose amplitude envelope you'd like to trace.**
- **Click on the Amplitude Envelope Mode icon.**

You are now in Amplitude Enveloping Mode. The waveform will appear gray, and a horizontal amplitude envelope line will appear near the top of the Y axis. The blank envelope line is now set to full on (attack).

- **Click on the Trace Envelope icon.**

The waveform's amplitude envelope will appear for editing. Most of Alchemy's normal editing functions will also process envelopes. Remember, envelopes are treated just like waveforms, and can be copied to the Clipboard for pasting into waveform windows, or superimposition as envelopes over other waveforms.

One of the most important concepts behind Alchemy's enveloping functions is the fact that there is no difference between envelope data and wavedata. In other words, any waveform can be copied to the Clipboard and pasted as an envelope, and any envelope can be copied to the Clipboard and pasted as wavedata. Using sound waveforms as amplitude or frequency envelopes can produce some very interesting results. And by pasting envelopes as wavedata, it is possible for you to maintain a library of different natural envelopes for future sound design tasks.

Amplitude Fit and Amplitude Scale



In order to alter the amplitude envelope of a waveform to match a newly created or pasted amplitude envelope, you will use either the Amplitude Fit or the Amplitude Scale commands (or icons).

Amplitude Fit automatically adjusts the amplitude of the waveform to fit the new envelope exactly – increasing certain amplitudes and decreasing others. Although this generally yields better results, it may cause distortion on low amplitude sounds, or those that contain areas of silence.

Amplitude Scale uses downward amplitude adjustment only – decreasing some amplitudes while leaving others unchanged. The effects of Amplitude Scaling are generally less dramatic, but will produce better results on low amplitude sounds, or those that contain areas of silence.


Looping Tools and Methods

Looping is the process of defining a section of a sampled waveform which is repetitive in nature and using that section to create an artificial “infinite sustain.” It is a subtle craft that relies on planning and listening. In this section you’ll find instructions for using Alchemy’s looping tools, and some looping examples. Undoubtedly you’ll develop your own methods and opinions as you become more familiar with Alchemy’s looping capabilities. For more information about looping, see the About Sound section of this manual, and the Looping Workshop in the Applications at the end of this section.

The Loop Cursors

The loop cursors are the most basic tools in Alchemy’s looping arsenal. They are simple vertical insertion points which are used to mark the loop start point and loop end point, and you are free to place them anywhere in a waveform. The loop cursors always stay where you put them until you move them to other spots, and their positions are always remembered when a sound file is saved.

Turning Loops On and Off

- Open any sound file.
- Make sure that the loops are not already on.
- Click on the loop cursors icon. 

You'll recognize the loop cursors as small solid vertical lines in the waveform window. If you are viewing your waveform windows with the axis markers on, or with the white space preference (from the Preferences... command on the Action menu), the loop cursors will also be marked with solid black triangular bases.

- Click and hold the Speaker icon, or hold down the space bar.

The current loop will play back as long as you hold down the mouse button or space bar.

Changing the Loop Points

- Place the mouse over any part of either loop cursor.
- When it changes to a left-right arrow, drag the loop cursor to a new position.
- Click and hold the mouse on the Speaker icon to hear the new loop.

This is one simple way to test possible loop areas. It can be useful to employ this method to try out different loop positions, but remember that all fine loop adjustments are probably best accomplished by using Loop Splice mode, which is explained on the next page.

Note: Some great looping shortcuts which can really simplify your looping work are the Select Loop and Loop Selection commands on the Edit menu.

Loop Splice Mode

Using Loop Splice Mode



As the name implies, Loop Splice mode is not just a tool. Along with Selection mode and Waveform Draw mode, it forms the group of three operating modes for the waveform display. Like the other modes, Loop Splice mode changes the waveform display so that it can be used for a specific purpose.

- **Choose the Full Zoom Out command on the Action menu to make sure that you're zoomed all the way out in the current waveform window**

If your loop is on, black vertical loop cursors should be visible in the waveform window.

- **If your loop is not turned on, turn it on by clicking on the loop cursors icon.**


The black vertical loop cursors appear in the waveform window.

- **Click the Loop Splice mode icon.**

Immediately you'll see the active waveform window change. Instead of the normal waveform display, you will now see a zoomed-in view of the loop splice point. On the left half of the display you see the waveform at loop end, and on the right you find the waveform at loop start. At the center of the display there is a dark vertical line. This line represents the splice point between loop end and loop start, and the loop splice display is the tool you'll use to make sure that the loop transition is a smooth and logical one.

- **Click on the scroll arrows near the bottom of the window to slowly shift either loop point.**

This moves the loop points one sample at a time, and is good for fine tuning.



- Click inside the gray regions of the scroll bar near the bottom of the window to automatically move the loop points to the next zero crossing.

Clicking on the left side of the scroll box moves towards the beginning of the sound file, while clicking on the right side moves towards the end of the sound. This automated zero crossing finder is a great looping tool, because zero crossings generally make good loop splice points.

- Click and drag either scroll box to make a large adjustment in the loop splice points.

Moving either scroll box to the left moves its loop point towards the beginning of the sound file, while moving either to the right moves towards the end of the sound.

- Click regularly on the Speaker icon or press the space bar to listen to your different loop splices.

A smooth transition between loop end and loop start (preferably on a zero crossing) will be quiet and click-free. If the loop is to be subtle, a smooth loop splice point is essential.

For some more hands-on experience, see the Looping Workshop in the Applications at the end of this section of the manual.


Crossfade Looping

- Select a waveform range which you'd like to loop.
- Choose the Loop Selection command on the Edit menu.

The selected range is now looped, but it's probably not a great loop.

- Choose the Xfade Loop... command on the Process menu.

The crossfade looping dialog box appears.



- Select the percentage of the crossfade by using the scroll arrows or box.

The percentage you choose tells Alchemy what percentage of the whole loop will be involved in the crossfade. If you set it to 100%, the whole loop will be part of the crossfade.

Frequency Analysis and Resynthesis

The Harmonic Spectrum Display

The default setting of 24% fades in the new sound over approximately the last quarter of the loop. This works well for most loops.

- **Click on the OK button.**

This executes your crossfade loop. You can click and hold on the Speaker icon to hear the new loop. You'll find it to be smoother than the previous loop, and click free. For some more information about looping, see the Looping Workshop at the end of this section of this chapter.

Alchemy is one of the few programs that not only analyzes the harmonic content of any sound, but lets you edit that harmonic content directly. With Alchemy you can look into the actual harmonic makeup of any waveform and adjust it exactly to your liking. You can remove or add whole frequency components, and perform precise high-pass, low-pass and notch filtering with a few mouse clicks.

All of this is possible due to the Fast Fourier Transform, which is an analysis and resynthesis algorithm allowing quick translation between the time and frequency domains. Alchemy's harmonic analysis and resynthesis hinges on two tools: The first is a separate window display which is used to show the harmonic spectrum of selected waveform ranges of up to 32,768 samples. Using this harmonic spectrum window you can see *and edit* each individual frequency band.

The harmonic spectrum display is the heart of Alchemy's frequency analysis and editing tools. Although it is a window, like the waveform display, its design and function are much different. It always appears as a separate window detailing the exact harmonic content of the selected waveform range.


Analyzing a Selected Waveform

- Select any waveform range of up to 32,768 samples.

- Click on the Analyze icon.



When you do this, a harmonic spectrum display window will open near the bottom of the screen. The harmonic spectrum display contains information which indicates the exact harmonic content of the selected waveform range. Harmonic information is shown as a number of vertical bars, each with its own amplitude. These bars are displayed on an X-Y axis which lists a band number below each bar, and amplitude in decibels (dB) that corresponds to the height of each bar. Since it is possible for an analyzed sound to have over 16,000 frequency channels, you'll find scroll arrows and a scroll bar at the bottom of the window. This allows you to scroll the display to the right or left in order to view more frequency components.



You can picture the harmonic spectrum display like this: Each numbered vertical bar represents a sine wave which is a component of the analyzed waveform. The amplitude of each bar tells you how much of that particular sine wave is in the original waveform. By adjusting the amplitude of the different sine waves, you increase or decrease the amount of that frequency in the analyzed waveform.

Before you can change the spectrum of an analyzed waveform range, you'll need to learn how to select frequency channels for editing. Here are the methods you'll use to do so.

Selecting Spectra for Editing

Selecting an Individual Harmonic

- Click the mouse on the harmonic spectrum window to make sure that it's selected.
- Click on any frequency channel (bar) to select it.

A selected frequency channel has a small rectangle on its tip. To get all of the information about the last selected frequency band, look in the palette's numeric display. All pertinent information about the selected frequency band or bands is located in the palette's numeric display. When a frequency band is selected, the numeric display will tell you four things about it.

- 1) The selected band's channel number (Ch:). The number of frequency channels present in any harmonic spectrum is a function of the analyzed waveform's size and complexity.
- 2) The frequency in Herz of the selected band (Hz:). From this you will know exactly what frequency you're adjusting if you change the selected band's amplitude.
- 3) The amplitude (in dB) of the selected band. This allows you to gauge and edit the exact level of the selected frequency in the analyzed waveform.
- 4) The phase of the selected frequency band (Ph:). Although a sine wave's phase is not thought to be audible to the human ear, it can greatly affect the way in which waves react with each other.

Selecting a Group of Harmonics

- Drag the mouse across a range of harmonic spectrum channels.

As usual, black dots appear at the tips of the selected channels. The numeric display shows only the information about the last selected frequency channel.

Extending a Selection

- **Hold down the shift key and drag the mouse over any harmonic spectrum channel(s).**

Do this when you wish to add a selection to an already selected harmonic range. This is the only method which allows selection of a non-continuous group of frequency channels.

Editing Harmonic Spectra

Once you have selected a frequency channel (or range) in the harmonic spectrum for editing, there are a number of processes which you may perform on that selection. Remember that no edits performed on a selected waveform range's harmonic information will be audible until the waveform is resynthesized.

Editing the Amplitude of a Single Frequency

- **Click and hold the mouse on the frequency channel's black selection box.**

This box is located at the tip of a selected frequency channel.

- **Drag the selected channel to a new amplitude.**

You can view the amplitude information of the selected channel as you edit it by looking in the palette's numeric display. Only one frequency channel at a time may have its amplitude changed. To reconstruct the new waveform to reflect your harmonic edits, you'll need to resynthesize it as explained later in this chapter.

Cutting Frequency Channels

- **Select a channel or channels by dragging the mouse cursor across the harmonic spectrum display.**
- **Choose the Cut command from the Edit menu.**

Cutting frequency channels removes them from the harmonic spectrum display and places them on the Clipboard for later pasting or mixing.

Removing a frequency band will totally erase that frequency from the waveform when you resynthesize. Removing a group of frequencies will erase that frequency range from the waveform when you resynthesize. So you can see that cutting frequency channels and resynthesizing is like running a sound through an extremely accurate notch filter. By increasing the number of frequency channels you cut, you are increasing the width of the notch.

Copying Frequency Channels

- **Select a channel or channels by dragging the mouse cursor across the harmonic spectrum display.**
- **Choose the Copy command from the Edit menu.**

Copying places a copy of the selected frequency channel or channels on the Clipboard, but it leaves the originally selected frequency channels in the harmonic spectrum display. The Clipboard channels can then be pasted or mixed into any frequency range.

Clearing Frequency Channels

- **Select a channel or channels by dragging the mouse cursor across the harmonic spectrum display.**
- **Choose the Clear command from the Edit menu.**

Pasting Frequency Channels

Clearing frequency channels erases them entirely from the harmonic spectrum display. No copy is placed on the Clipboard for later pasting or mixing, and they can not be retrieved. Clearing a frequency band will totally erase that frequency from the waveform when you resynthesize. Clearing a group of frequencies will erase that frequency range from the waveform when you resynthesize. So you can see that clearing frequency channels and resynthesizing is like running a sound through an extremely accurate notch filter. By increasing the number of frequency channels you clear, you are increasing the width of the notch.

To clear all frequencies above a selected channel:

- **Choose the Clear Above command from the Edit menu.**

All frequencies above the selected range will be removed entirely from the harmonic spectrum display. This is basically a very accurate low-pass filter which removes all frequencies above the selected channel when you resynthesize. The Cut Above command is only available when a single frequency channel is selected.

To clear all frequencies below a selected channel:

- **Choose the Clear below command from the Edit menu.**

With this command, all frequencies below the selected range will be removed entirely from the harmonic spectrum display. This is basically a very accurate high-pass filter which removes all frequencies beneath the selected channel when you resynthesize. The Cut Below command is only available when a single frequency channel is selected.

- **Make sure that you have cut or copied at least one frequency channel.**

Mixing Frequency Channels

All amplitude and phase information about the cut or copied bands is saved, but you are free to change their frequency by pasting them somewhere else in the harmonic spectrum.

- **Select a channel or channels as the destination range by dragging the mouse cursor over them.**
- **Choose the Paste command on the Edit menu.**

When you choose the Paste command, the frequency channels on the Clipboard will replace the selected frequency channels. Only the selected paste channels will be replaced with the Clipboard channels. If you have ten channels on the Clipboard, and only one paste channel is selected, only the left-most Clipboard channel will be pasted. If you select a paste range that is wider than the Clipboard range, the contents of the Clipboard will only be pasted over the beginning of the paste range. The end of the paste range will remain untouched.

- **Make sure that you have cut or copied at least one frequency channel.**

All amplitude and phase information about the cut or copied bands is saved, but you are free to add their channel characteristics elsewhere by mixing them with another range in the harmonic spectrum.

- **Select a channel or channels as the destination range by dragging the mouse cursor over them.**
- **Choose the Mix command on the Edit menu.**

When you choose the Mix command, the Clipboard channels' amplitudes and phases are taken into consideration and are added to the selected destination channels. By cutting frequency channels you want to remove, and mixing them elsewhere in the harmonic spectrum, you can retain the overall amplitude of the waveform after resynthesis.

"Out-point" Editing

Resynthesizing the Original Waveform

Simply clearing frequency channels subtracts power from a waveform, and thereby decreases amplitude. Try experimenting with frequency band pasting and mixing to see exactly what effect this has on the waveform.

To Paste or Mix the Clipboard starting at range end:

- **Select a destination frequency range.**
- **Hold down the option key while you choose the Paste or Mix command on the Edit menu.**

This version of the Paste and Mix commands takes the frequency channel range which is currently on the Clipboard and pastes or mixes it into the newly selected harmonic range beginning with range end and Clipboard end. This is the only way to match up the high end or "out" point of a paste or mix.

- **Click on the Resynthesize icon.**



The Fast Fourier Transform is immediately reversed in order to reshape your original waveform to reflect the harmonic changes you've made. Like analysis, this may take a few moments, but when the mouse cursor reappears you can click on the Speaker icon to listen to the new waveform. You may also want to look closely at the shape of the new waveform to see what your frequency edits have done. You are now free to reanalyze and/or resynthesize the same range to make more changes, or continue on with other tasks.

Digital Signal Processing Functions

Alchemy's digital signal processing tools may be divided into four realms: sample rate conversion, digital EQ, Time Scaling and Pitch Shifting, and DSP effects. Sample rate conversion lies at the heart of Alchemy's universality, because it allows you to adjust minutely any sound file's sample rate or number. Digital EQ brings a capability usually relegated to hardware boxes to the waveform editing world by offering quick and clean equalization of any size waveform. Time Scaling and Pitch Shifting are sophisticated functions that allow extremely accurate adjustments to be made to pitch (without altering duration), and duration (without altering pitch). And finally, DSP effects represent a new frontier waiting to be developed. In the near future they promise to free the sampling artist from many cumbersome hardware effects.

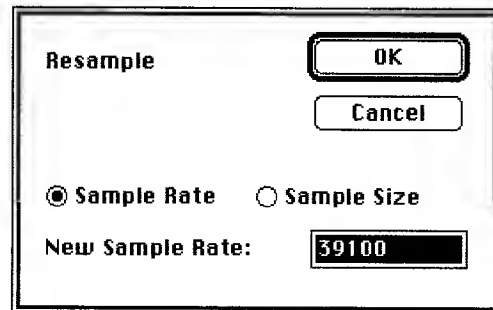
Sample Rate Conversion

One of Alchemy's most powerful DSP tools is its Resample function. Resampling is really the key to breaking the communications barrier between different sampling machines. By using this function, you can exactly adjust the sample rate and number of samples in any sound file. Since the conversion can be set in single-sample increments, any sampling format can be used or created. And since Alchemy uses a resampling function, rather than simple linear interpolation, you can actually increase a sound's sampling rate or number of samples with reasonable success. Here's how it works:

To Resample a Sound

- Open the sound file you wish to resample.
- Select the entire waveform by double clicking on the Selection mode icon.
- Choose the Resample command on the Process menu.

The resampling dialog box appears. In it you'll find two boxes. One lists your current sample rate, and the other the number of samples.



- Type in the sample rate (or number of samples) at which the sound should be resampled.
- Click on the OK button.

A dialog box appears asking you if the number of samples in the sound file should be changed to correspond to the new rate (or the rate to the new number).

- Click on the dialog's OK button.

After a few moments your waveform window reappears. The sound in it now has a different sample rate. To make sure, you can select the Soundfile Setup command on the File menu. The sample rate and number listed there are part of the sound file's characteristics. You can now send that sound with its new rate and number to any sampler on your network, which is particularly useful if you have any samplers with non-adjustable playback rates.



Important Note: Be careful not to down-sample a sound file (decrease its sample rate) and then save it over the original, higher-fidelity sound file. If you plan to save a down-sampled file, use the Save As... command to save it under a different name.

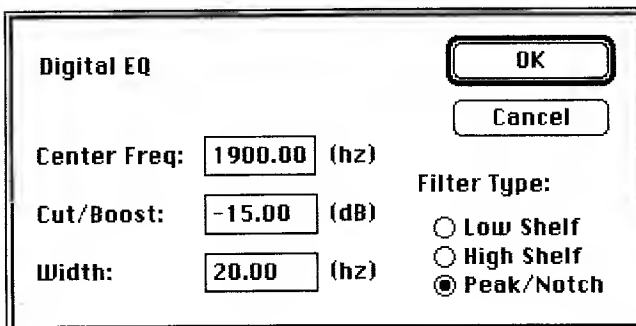
Digital EQ

Like harmonic spectrum editing, digital EQ directly affects the frequency content of a selected waveform range. Although Alchemy's digital EQ function may seem to overlap with that of the harmonic spectrum window, the two are actually quite different. When you construct a filter using the harmonic spectrum window, its characteristics are razor-sharp. Cutting a frequency of 1266 Hz touches *only* that frequency. The digital EQ function, on the other hand, is a more gentle-walled EQ, similar to high-quality hardware EQs. It affects the selected waveform range with a characteristically smoother curve. It also has no limit on the size of a waveform range it can equalize.

Digitally Equalizing a Waveform

- Select the waveform range that you want to equalize.
- Choose the EQ... command on the Process menu.

The digital EQ dialog box appears.



The screenshot shows the 'Digital EQ' dialog box. It has a title bar 'Digital EQ' and two buttons: 'OK' and 'Cancel'. The dialog contains three input fields: 'Center Freq:' with the value '1900.00' and unit '(hz)', 'Cut/Boost:' with the value '-15.00' and unit '(dB)', and 'Width:' with the value '20.00' and unit '(hz)'. To the right of these fields is a section titled 'Filter Type:' with three radio button options: 'Low Shelf', 'High Shelf', and 'Peak/Notch'. The 'Peak/Notch' option is selected with a filled circle.

- Type in the center frequency which you wish to cut or boost.
- Type in the amount (in dB) by which you wish to cut (-) or boost (+) the above frequency.

Time Scaling and Pitch Shifting

Use a minus sign before the number to indicate a cut, and no sign (or a plus sign) to indicate a boost in the listed frequency.

If you're building a notch filter:

- **Choose the frequency width in Herz of the cut or boost.**

Use the width to define how wide the notch will be. A width of 20 Hz, for example, will affect all frequencies within 10 Hz of the center frequency.

- **Click on the button in front of the filter type you desire.**

A high shelf will boost or attenuate all of the frequencies below the cutoff frequency, and a low shelf will boost or attenuate all of the frequencies above the cutoff frequency. A notch filter simply turns down a range in the sound, and the width of that range is dictated by what you type into the Width box.

- **Click on the OK button.**

If you are equalizing a long sound, it may take a moment before the waveform window appears again, but when it does, your EQ is complete.

Until now, sound designers have traditionally been limited in the types of adjustments that they could make to the pitch and duration of sampled sounds. It has always been relatively simple to change the pitch of a sound, but such changes required that the duration of the sound also be adjusted. Similarly, the duration of a sampled sound has generally been adjustable only if the pitch of the sound could be changed.

By using new and sophisticated algorithms, Alchemy has broken ground on a new sound design realm. As you have seen, Alchemy's Frequency Enveloping mode is one way to alter the pitch of a sound over time. However, Alchemy also

allows both the time and pitch of any sampled sound to be adjusted independently, and very accurately. Even outside of synthesis-oriented sound design, the implications of this are far-reaching. For example, duration-independent pitch shifting makes it possible to create multisamples from a single source sample, and to create complex digital effects such as harmonizing. Pitch-independent time scaling allows sound designers and post production engineers to adjust the duration of sounds to precise time (or SMPTE frame) lengths.

Time Scaling a Sound

- Select the range to be scaled, or select the entire sound.
- Choose the Time Scale... command on the Process menu.

The Time Scale dialog appears.

Time Scale		OK
		Cancel
Time Scale Factor:	1.3000	Calculate
End Point: 0.2762	0.5381	(seconds)
Duration : 0.2062	0.2681	(seconds)

- Type in the new end time for the range or a new duration for the range, or type in a scale factor describing the relationship between the scaled file and the original file.

You may type a value into any of these boxes, and click on the Calculate button to see the other matching values. For example, enter a scale factor of 1.2 on a 10 second sound and click on the Calculate button. The duration value is



Pitch Shifting a Sound

automatically updated to show that the scale factor would produce a sound with a duration of 12 seconds. Durations and end times are always indicated in the units you have chosen using the Axis Units pop-up menu on the Windows menu.


- Click on **OK** to execute the time scale.

A thermometer-style indicator will appear and indicate the time scale progress. After processing is complete, the waveform window will reappear and the sound will have its new duration.

Note: If you are attempting to retain the original timbre of the sampled sound, time scaling produces its best effects when used with scale factors between 0.7 and 1.3. More extreme factors may alter the tone of the sound, producing effects that may be of interest to sound designers.

- Select the range you wish to pitch shift, or select the entire sound.
- Choose the **Pitch Shift...** command on the **Process** menu.

The Pitch Shift keyboard dialog appears. A dark gray key indicates the selected sound's current pitch, and the key is labeled at the bottom of the dialog.



- Click on the keyboard key that represents the new pitch you wish the sound to have.

Digital Signal Processing Effects

The 'Transpose by' box indicates the size of the planned transposition. **Note:** You can also set the size of the pitch shift by typing a positive or negative value into the 'Transpose by' box.

- **To make sure the pitch shifted sound's duration matches the original duration, click on the box in front of "Preserve Duration."**
- **When all of the pitch shift settings are to your liking, click on the OK button to execute the pitch shift.**

A thermometer-style indicator will appear and indicate the progress of the pitch shift (and time scale, if applicable). After processing is complete, the waveform window will reappear and the sound will have its new pitch.

In many instances DSP effects offer a truly viable alternative to hardware effects. A DSP effect is an algorithm or method which creates its effect by mathematically altering the original waveform to make the processed waveform. You could, for example, create a DSP 'spring reverb' algorithm which would take any waveform and mathematically process it to become a 'spring-reverbed' version of itself. Since all of the actual processing takes place as a manipulation of numbers in the digital realm, no line noise or distortion is added to the signal, and the new waveform comes out as clean as the original.

Digital signal processing effects are to the sampling world what 'hardware' effects are to the analog world. Over the next page you'll find a simple example of how you can create your own DSP effects using Alchemy as your workshop.

Example: Building a Simple Echo

- Open a sound file which contains a single short, relatively sharp sound.
- Select the entire waveform by choosing the **Select All** command on the **Edit** menu.
- Store the current view in view memory 1 by clicking on the palette's **SAVE** button, and then on the number 1.
- Choose the **Soundfile Setup** command on the **File** menu.

When the dialog box appears, look at the number next to **Sample Size**. We'll assume that you want to build an echo that repeats three times. Use the Macintosh calculator to multiply the **Sample Size** number by four. This new number describes how many samples you'll need to create your echo sound file. You multiplied by four because your end product will consist of the original waveform followed by three copies.

- Enter the new number in the **Sample Size** box in the dialog.
- Click on the **OK** button.
- Hold down the **Command (⌘)** key and click on the **Zoom Out** icon.

Your whole sound file is now visible. You'll notice that the sound file is much longer than before, but the end is empty.

- Click view memory one.
- Choose the **Copy** command on the **Edit** menu
- Double click on the **Selection** icon to select the entire waveform.
- Click on the **Replicate** icon.

There will now be four copies of your sound, one after the other.

- **Click the Fade Out icon.**

Now click the Speaker icon to hear your creation again. Although this is a very basic effect, it should give you some idea of what you can do simply by manipulating multiple copies of a sound.

With the information you have now read, you should be able to use Alchemy to accomplish your sound design goals. If you'd like to look at some concrete examples of how Alchemy can be used, read through the Applications which follow.

Sample Applications

This section of the Alchemy manual contains some practical applications which should help you understand how to use your distributed audio network. Although the first four are general, the last one is for SP-1200 owners only. These examples are designed to take you a bit deeper into the realities of using Alchemy, and hopefully they will answer some important questions and speed your transition into the network environment.

Application 1

Looping Workshop

In both the About Sound and Using Alchemy sections of this manual you have read some basic information about loops. In this application we'll look at some more in-depth information and some actual techniques you can use to construct and fix such loops. We'll start out with general background and move through a simple spliced short loop, correcting a volume bump, and the more complex mirror and crossfade loops. After many applicable topics you'll find a Do it Yourself explanation which uses a sound file supplied on your Alchemy sound disk to give you a concrete demonstration of the concepts.

Short Loops

Short loops are usually loops of 16 wave cycles or less, and are generally best suited for sounds which have little or no harmonic development. A sawtooth or square wave would be a perfect waveform for a short loop. A sampled buzzer or flute probably would also. A horn section or human voice would probably not be well suited for a short loop.

Use a short loop any time you have a waveform which seems to repeat regularly with little or no change in wave shape. However, when you are dealing with a waveform that has extreme changes in wave shape over time, your best bet may be a long loop.

Long Loops

Most complex sounds, especially natural sounds, have waveforms which change shape radically over time. This is a reflection of the many harmonic changes taking place as these sounds evolve. With sounds such as these, short loops are usually useless. Since there is no single wave shape or group of wave shapes which can be looped over and over, another method must be found.

Looping Techniques

This other method is the long loop. Instead of looping a single wave cycle, a long loop loops a major portion of the sampled sound. One way to think of it is to picture a two-second sampled sound of your voice singing 'laaa...' A long loop can be constructed to include the entire '-aaa...' section of the sound. Since a loop this large may contain hundreds or thousands of wave cycles, it often doesn't lose the harmonic qualities of the original sound. In fact, sometimes a loop of this nature can be virtually unnoticeable.

Long loops are really a fertile creative ground for instrumental sound design, and there are many methods and tricks for creating them. Later in this application you'll find discussions of some tools and methods for creating long loops. This is new territory, so you won't find every question answered. What you will find are some opportunities to develop methods of your own.

Let's assume you've sampled your voice singing the word "laa..." Using a visual editing system, you are now able to view the entire waveform from start to finish. You are also capable of isolating any section of the sound for listening or editing, and you can set loop start and end points at will.

Given these assumptions, you have many options for short or long loops. You could place the loop start point at the beginning of the sound and the loop end at the end of the sound. This would make the whole sound repeat over and over, but that's an obvious loop, and it's not very natural. As a final assumption, then, we'll assume that you want to make the sound as natural as possible by looping only the "aa..." section of the sound. This will leave a "laa..." sound that starts only once and sustains indefinitely.


Depending on the nature of your sampled sound, you might build a short loop (of a few cycles) or a long loop (of many cycles). Using your visual editing system, you would isolate and listen to a few sections of the "aa..." sound and look to see if you can find a waveform range which seems to repeat. You

A Click in the Loop

would then place the loop start on (or just after) the very beginning of the isolated section, and place the loop end at (or just before) the very end of the section.


If you are an extremely lucky person, you may be done. You'll know this is the case if you hear your voice smoothly sing "laa..." for as long as a key is held down. The chances of this are very low, so don't worry if your loop sounds funny.

The problem you will run into most often is a click or bump in the loop. Since this is generally caused by incorrect positioning of the loop start and end points, it is usually easy to fix. To visualize how to do it, picture the waveform at loop start and loop end.



As you can see from this diagram, the waveform at loop end does not match the waveform at loop start. Since the transition is not smooth, there's a nasty click when the loop begins. To fix this, you will need to adjust very slightly either the loop end point or the loop start point. For this illustration let's adjust the loop end point.

You can see that the loop end waveform would match perfectly with the loop start if the loop end splice point were moved a tiny bit to the left. Good visual editing systems allow you to slide the loop end point over until the waveforms match up. This gives you the impression that you are pushing the entire loop end waveform to the right, although you're not really moving the waveform, just the loop end point. Do this until the transition is smooth.



This type of correction is very simple, but unfortunately a smooth waveform transition does not always guarantee a good loop. You may need to try a few looping points before you get something you like.

There are some problems, however, which you can't fix by adjusting the loop start and end points. If the volume or tone of the looped interval changes over time, the splice point may be perfect, but the loop may seem to surge or swell repetitively. This often draws unwanted attention to the loop and makes it sound unnatural. To fix these problems you'll need some more advanced tools.

**Try it
Yourself:
Correct a Loop
Click**

**Volume
Changes
in the Loop**

- Open the sound file entitled “Loop Sine” on the Alchemy sound disk.
- Click and hold on the Speaker icon (or hold down the space bar) to hear the loop.

Although this sound file contains a sine wave, this is not a sine tone because the waveform is looped incorrectly. Since a sine tone consists of only one fundamental tone, it is harmonically pure and very quiet. The waveform you just heard is far too loud and edgy. You will now use Loop Splice mode to loop this sine wave correctly.

- Click on the Loop Splice mode icon.
- Click and hold on the farthest left scroll arrow at the bottom of the window.

As you do this you will see the loop end waveform slide to the right as the loop end point is moved closer to the beginning of the sound.


- Hold down the space bar to hear some successive loops.

As the waveforms come closer to matching up, the tone will gradually rise, and as the waveforms get close to matching it will become quieter and more ‘sine-like.’ When the transition is smooth your job is done. This is the basic procedure you’ll follow to fine tune any loop transition.

Sometimes you can’t avoid having a volume change in the waveform section you wish to loop. If you had moved the microphone away from your mouth at the end of your “laa...” sample, the decrease in volume over time would probably have left a volume bump in your loop territory. Even after you created a smooth loop splice, this bump would repeat with every loop. To correct this problem you can adjust the volume of a waveform section by using a *scaling* function.

Scaling is the act of defining the percentage of volume increase or decrease you want over a selected waveform section. Scaling is usually executed so that all volume levels in the selected waveform area are decreased by the same factor.

For the sake of experimentation, let's assume that the volume bump is at the very beginning of the "laa..." loop.



When only the beginning of the waveform range is selected for scaling, only the beginning amplitudes are decreased. This has the effect of turning down the beginning volumes while leaving the ending volumes the same, which is exactly what we wanted to do. As long as this waveform section is looped smoothly after scaling, it should have little or no volume swell. *Hint:* In order to avoid creating any waveform irregularities when you scale, always make sure that the scale range begins and ends on a zero crossing. Since scaling a zero leaves it unchanged, the transition between scaled and unscaled waveforms will always be smooth.

Tone Changes in the Loop

This brings us to the next, and most complex looping problem. What do you do when there is a noticeable change in tone during your selected looping interval? For example, what if the beginning of the “laa...” sound is brighter and more brassy than the end? If you think back to what this means, you’ll see why this is a more serious problem.

The harmonic content of any waveform determines its shape. Since “tone” is really the same as harmonic content, two waveforms with different tones would have different shapes. When you build your loop, you will be trying to engineer a smooth, seamless transition between very different wave shapes, which is pretty much impossible. Even if the splice point matches perfectly, the tone will jump radically, and will make you very conscious of the loop.

The best way to deal with extreme tonal differences within a loop is find more stable source material. Regardless of what you do, there will be a tonal surging or swelling. You’d be better off choosing a smaller loop area with less (or no) tonal variation. Still, even minimal tonal change within a loop will require some thoughtful action. To make a loop like this work, you have to find some way of making the shapes of the loop start and loop end waveforms as similar as possible. At the same time, you have to make sure that all tonal changes within the loop are slow and evolving.


We’re going to look at three ways to build loops with smooth tone evolution: The mirror loop, the mirror crossfade loop, and the standard crossfade loop. The mirror loops are more cumbersome and need to be built by hand, so they are only explained conceptually. If you want to try them, open your own sound and follow along. Automated crossfade looping is one of Alchemy’s features, so a sound file has been included on your sound disk to demonstrate the standard crossfade loop. Although the standard crossfade loop is currently the most popular, the other methods discussed here may give you some good ideas.

Mirror Splice

The mirror splice method is based on the concept that the harmonic content of any waveform is the same backwards or forwards. Although this is a usable concept for many sounds, for others it doesn't work well at all. To get a good feel for this, try it out a few times.


Executing a mirror splice takes a bit of planning. The idea is to make a copy of the first half of the loop, reverse the copy, and make it the second half of the loop. This way you can be sure that the wave shape at the end of the loop will match the wave shape at the beginning of the loop (although it will be backwards). Here's an example:

A Mirror Splice Loop



By doing this you manage to make the loop end wave shape match the loop start wave shape, but there's still something wrong. The wave shape at the loop midpoint looks irregular, and will probably cause a click when the loop is played. To fix this, there is one more step we have to go through. This step is based on a mirror splicing law: The slope of the waveform must remain the same on both sides of the mirror point (the loop midpoint). Although this sounds confusing, there's an easy way to do it.

Finishing the Mirror Splice Loop



By inverting (turning upside-down) the new mirror section of the waveform, the loop's midpoint problems are taken care of, and the waveform's slope remains continuous on both sides of the mirror point. You can do this, because turning a waveform upside-down has little or no effect on the sound.

You should also note that it's wise to have your loop midpoint fall in a place where the waveform crosses zero. This way the original waveform and its inverted mirror will always match up perfectly.

Once the mirror splicing is finished, make sure that the loop end is spliced smoothly with the loop start, and you're done. This will always be easy, since the loop end waveform is upside down, but identical to the loop start waveform.

Although this method looks great on paper, it doesn't always work. Natural sounds are so complex that any change in their harmonic evolution will be noticeable. The mirror splice used here is a hard-edged splice, so the loop midpoint will be a radical transition by nature. A good mirror splice loop will have no clicks, but may sound less smooth and more artificial than the original source sound. A bad mirror splice loop will be very obvious, and may contain clicks and bumps.

Splice Loops vs. Crossfade Loops

Obviously, creating long loops presents some new challenges. Creating a short loop was easy because you wanted every successive waveform to sound the same. When you create a long loop, you are asking for a lot.

The potential problem here is that there is often quite a bit of waveform between loop start and loop end. If the shape, and hence the harmonic content of the waveform changes greatly anywhere in the loop, there will be a very noticeable timbral 'bump' every time the loop splice point is passed. Even if you line up all splice points perfectly, any radical change in wave shape will make your loop very noticeable. With a straightforward splice there may be no way to fix it.

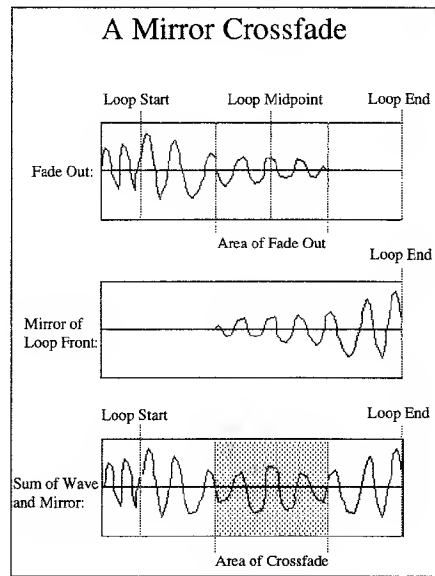
This is where the crossfading comes in. Although it's not a cure-all, crossfading can often be the salvation of a bad situation. The concept behind a crossfade is simple. If you're looking for the best transition between one wave shape and another, simply mix both wave shapes together over a range, fading one smoothly out as the other is faded smoothly in.

Essentially, that is the crossfade. There are, however, a few things you'll need to pay attention to. You have to make sure that a crossfade itself happens over a large enough area not to be intrusive. If the sound fades too quickly from one sound to another, you'll still have a timbral 'bump' in your final product. If the sound fades from one sound to another over too large a range, you run the risk of requiring a massive loop area.

Mirror Crossfade

In practice, all crossfades are accomplished by adding the waveform section that is being faded out to the waveform section that is being faded in. In a mirror crossfade, the waveform section to be looped is faded out across the loop midpoint. The center of the fade will always be the loop midpoint.

Once you have finished the fade out, make a copy of the whole loop, from the loop start to the point where the fade out is complete. This copy is then reversed, and often it will help if you invert it to make the loop start and loop end waveforms match correctly.




When all of this is complete, you'll add the reversed copy to the original. Make sure to match up the beginning of the copy's fade in with the beginning of the original's fade out. This assures you that the overall volume of the loop center will remain consistent.

Fade Slopes

If the volume in the loop always seems to drop after the waves have been added, you may need to try a different fade slope. The simplest and most common fade is a linear fade. In a linear fade, amplitude is decreased in a straight line. Halfway between fade start and fade end the amplitude has been decreased to 1/2 of its former value.

Another, often better option is the 3dB fade. It fades using a gentle slope, rather than a straight line. A 3dB fade decreases amplitude at the fade halfway mark to 7/10 of its former value (compared to 5/10 in the linear fade). By using a 3dB crossfade, you ensure that the amplitude at the center of the crossfade area will be 4/10 higher than with a linear crossfade ($7/10 + 7/10$ versus $5/10 + 5/10$).




It's generally wise to use 3dB crossfades with more complex sounds. The linear fade is useful primarily for simple waveforms or single instrument samples, which have less waveform development and are more regular.

By now you've finished your mirror crossfade loop. You'll find that the center of the loop is much more stable, and less likely to click, pop, or bump. As long as you are careful to line up your loop end/loop start splice point, the end of the loop should present no problems. Still, there's no guarantee that your loop will be what you wanted. Perhaps the whole mirror section of the loop sounds unnatural. After all, it is backwards, and probably upside-down. There is another, perhaps more natural way to deal with a tonal change in your loop.

Loop Start/ Loop End Crossfade

You've seen how useful crossfading can be. It allows you to avoid all of the problems of hard-edged splicing, and is only slightly more difficult. One of the most versatile crossfade methods is the loop start/loop end crossfade, which is often just referred to as a "crossfade loop." If you have a loop that you like, but you can't get rid of the loop splice pop and slight timbral difference between loop end and loop start, this is definitely worth trying.

The loop start/loop end crossfade is probably the most popular crossfade loop method. It is based on this very simple idea: The best transition into the loop start waveform is its natural transition. The section of waveform which usually comes right before the loop start is, by nature, the perfect transition into it. Unfortunately, when you establish a loop this transition is never repeated. Your sampling machine will play the natural progression of the sound from its very beginning to loop end. Then playback begins again at the loop start point. The wave section immediately before the loop start point, the perfect transition, is only played once.



In a loop start/loop end crossfade, the section of waveform immediately before the loop start is crossfaded with the section of waveform at the loop end. As the loop nears its end, the loop's own waveform is faded out, and a copy of the waveform from immediately before the loop start is faded in.

Since we know that this copy forms a perfect transition into the loop start, we rid ourselves both of timbral problems, and loop splice noise at the splice point. Samplers or programs with automated crossfade looping use this exact process to smooth the loop splice transition.


To execute a loop like this, you simply need to fade out the waveform at the end of the loop. The fade zone can be any length, but make sure you know its exact size and fade curve. Next, copy a section of the waveform which comes immediately before the loop start point. It must be at least as long as the section of the loop that you just faded out. Take the copy and fade in the beginning. The length of the fade in, and its fade curve, should be exactly that same as the loop end fade out you made earlier.

Finally, add the faded-in wave copy to the end of your loop, making sure that the fade-in area is perfectly superimposed over the faded-out area. To finish your loop, adjust the loop splice point so the transition is smooth. The loop should sound natural and click-free, because the loop splice point is a copy of the wave's natural harmonic evolution. If the loop evolves too quickly and produces a rapidly repeating swell, try using a longer crossfade area. If the change is too slow and undynamic, try a shorter crossfade area.

There is one final concern with certain crossfade loops. If your loop end point comes before the actual end of your waveform, you may have another step. By creating the crossfade loop, you made sure that the loop end and loop start wave shapes match, and your loop sounds smooth. Unfortunately, you changed the wave shape at the loop end to accomplish this. If you were now to trigger the sound a single time, with no loop, you'd find that the loop end waveform no longer blends smoothly into the waveform after it. This would happen every time you let go of a key, so it would be pretty noticeable.

To fix this you need to execute another crossfade. By using the same principles explained above, you can copy a section of waveform from the loop start range, because this now flows naturally out of the loop end. Then crossfade the copied range

with the waveform after loop end, being careful to fade out the copied range as you fade in the waveform end. This will leave you with a finished, click-free loop with a smooth transition into the original waveform's natural end. Here's an illustration:



From reading though the above explanation you've probably realized what an incredible pain it would be to do all of this every time you want to create a crossfade loop. In truth, any good visual waveform editor will have an automated crossfade looping function which lets you decide the size of your crossfade range, and then does the rest for you. Although all crossfading could be done step-by-step, automation can make it both fast and easy.

Other Types of Loops

Although we've been concentrating on one type of loop, there are others. Our loops have all been 'forward only' loops, which means that playback goes from loop start to loop end, and then jumps back in at loop start. There are also 'backward/forward loops, which repeatedly play from loop start to loop end, and then turn around and play from loop end to loop start. This electronically emulates the concept of a mirror loop, and

thereby avoids some of the problems associated with tonal changes in forward only loops. However, like a mirror loop, it gives you two splice points to worry about.

There is also another wholly different loop type. We built our loops, like the “laa...” loop, to repeat for as long as a key is held down on the keyboard. These kinds of loops are called “sustain loops,” and most samplers only allow this type of looping. Some samplers, however, also have what is called a ‘release loop.’ This allows you to define a second loop. The sustain loop repeats for as long as a key is held down, and the release loop repeats during the sound’s defined release period (the period after you let go of the key). The release loop offers even more sampling flexibility by allowing digital emulation of effects like echo, and it is only one of a number of new ideas. Since technology doesn’t seem to be slowing down, it would be safe to say that we’ve barely begun to investigate the different types of loops and looping methods.

You’ve probably noticed by now that looping is a strange, almost mystical practice. The mixture of predictability and unpredictability is inspiring and frustrating. You could probably spend your life trying to complete a single smooth loop if it was nasty - and that’s the point. Choose your raw material carefully. Some loops will never sound the way you want them to. Extreme harmonic changes in a wave tell you that a sound is developing. A loop restricts that motion, and will probably never sound right. Loops are a subtle creative tool for increasing a sound’s flexibility and conserving memory. They take careful planning, and they seldom work miracles. However, if you don’t push them too far, they can be very inspiring.

Application 2

Building and Using a Universal Sample Library

Alchemy's sound environment has been specifically designed to make possible the creation and use of a central sound library. Ideally this central library would contain 16-bit sounds (many of them stereo) which are sampled at the highest possible rate. High sampling rates give you higher fidelity over a broader frequency range. And since Alchemy's Audio IFF files remember key range, sounds which constitute an entire keyboard would be stored in a folder together (a full keyboard of piano samples, for example). This folder full of high-fidelity sounds could then be opened and sent to a network sampler without worrying about assigning key ranges.

However, when you are dealing with a central sound library sampled at 44.1 kHz, for example, you may need to make some sample rate adjustments when you want to send your sounds to a network sampler. After you decide on a destination sampler for your sounds, you will need to pick the sample rate you'll be using on that sampler. Then, before you send your library sounds to that sampler, you'll need to resample the sounds to match the sampler's rate. If you can set your destination sampler to play back at the rate your library sounds were sampled then you avoid this step, but sometimes this will be impossible or undesirable.

This application presents an example of the resampling and sending procedure you might follow to transfer a full keyboard of sounds from your library to a network sampler. For the purposes of demonstration we'll assume that your library sounds are sampled at 44.1 kHz and your sampler is set to a rate of 28 kHz. To go through this application, set your destination sampler to a sample rate of 28 kHz and start up Alchemy.

- **Make sure your destination sampler has been added to your network.**

This process is explained under DAN (Distributed Audio Network) near the beginning of the Using Alchemy section of this manual.



- **Choose the Open Group... command on the File menu.**

The standard open dialog box appears. If you're running from your program disk, insert the diskette with your high-rate sounds on the other drive. The sound file titles appear in the Open box.

If you're running from a hard disk:

- **Locate and click on the first of the multi-sampled sound files, then click on the Open button.**

The first sound file appears in an open waveform window on Alchemy's desktop. Then the open dialog box reappears.

- **Repeat this procedure until all of the sound files for the full keyboard range are open.**

Each waveform appears on the desktop in a waveform window stacked on top of the last opened sound. When all of the files are open, the open dialog box appears again.

Note: If you do not have enough memory to open all of these files simultaneously, you will need to complete this procedure by working on one sound file at a time.

- **Click on the dialog box's Cancel button.**

The open dialog disappears and you are back on the desktop where all of your open files are stacked.

- **Choose the Tile command on the Windows menu.**

The open sound files are now evenly spaced and sized so you can see all of them.

- **Activate the first waveform window by clicking anywhere inside of it.**

Its title bar is now highlighted, and it's ready for editing. You can click on the Speaker icon if you want to hear what it sounds like.

- **Choose the Soundfile Setup... command on the File menu.**

The Soundfile Setup dialog box appears. It contains all of the vital information about the currently selected sound. Notice the high sample rate, which verifies that the sound was sampled at a rate above the sample rate of your eventual destination sampler.

- **Click on the Cancel button to close the dialog box.**
- **Choose the Resample... command on the Process menu.**

The resampling dialog box appears. In it you find two buttons and a highlighted box. The buttons can be used to select the sample rate or sample number of the current sound. The box lists the current sample rate (or number) and you'll use it to enter the new rate (or number).

- **Click on the button in front of Sample Rate.**

The sample rate of the current sound file appears in the highlighted box.

- **Type "28000" in the highlighted New Sample Rate box.**

- **Click on the OK button.**

A dialog box appears and asks if you want to decrease the number of samples in the current sound file. You will need to decrease the number of samples in the file if you want to resample it to a lower rate.

- **Click on the OK button to decrease the number of samples.**

The actual resampling is now taking place. When it is done, the waveform window will reappear with the resampled waveform.

- **Choose the Soundfile Setup... command on the File menu.**

The Soundfile Setup dialog box appears. Notice that the new sample rate is 28000, which verifies that the sound was resampled correctly.

- **Click on the Cancel button to close the dialog box.**
- **Click on the Speaker icon to hear the sound.**

The only thing that has changed about the sound is its sample rate. Its pitch remains unaltered. Since you decreased the sound's sample rate, you have lost some fidelity, but this should be minimal.

- **Choose the name of your destination sampler at the bottom of the Network menu.**

A check mark appears in front of it.

- **Choose the Send Sound command on the Network menu.**

After a moment the keyboard dialog appears.

- **Use the incremental arrows next to "Instr:" to select the destination instrument (preset).**

This is only available on samplers which have multiple instruments/presets (EPS, Emax, Akai S-900).

- **Click on the Assign New Voice box.**

An "X" appears in it, and the key range and Unity key of the sound you're about to send is automatically selected. This is the key range and Unity key information which goes with the original sound. This is stored as part of any Audio IFF file.

- **Click on the OK button.**

If your sampler currently has any sounds in memory, dialog boxes will appear and ask if it's OK to delete those sounds to make room for your new sounds. Clicking on the OK button only clears room in memory, it doesn't delete any sounds from your disk.

- **Click on the OK buttons in any delete voice dialog boxes which appear.**

You will only be deleting voices from your samplers memory, not from any disk. When you have completed this, the transfer begins. A thermometer-like display appears and lets you monitor its progress. When the transfer is complete, the Alchemy desktop and waveform windows reappears. The original source sound has now been resampled and sent to its correct key range on your destination sampler.

- **Choose the Close command on the File menu.**

This tells Alchemy that you want to close the sound file you just resampled and sent. A dialog box appears and asks if you want to save the changes you've made to the file.

- **Click on the dialog box's No button.**

This is very important. You don't want to save the changes you made, because they involved the down-sampling of your source sound. Since you decreased the sample rate of the original sound, you also decreased its fidelity. If you save the changes you made, you will only have a 28 kHz source sound instead of the original 44.1 kHz sound.

- **Repeat the resampling and sending procedure for each sound on your Alchemy desktop.**

Make sure that you close each file without saving the changes after you have completed your transfer. Otherwise you'll replace some of your original source sounds with lower quality versions.

When you are done, you will have a full sampler key range with all of the source sounds in their correct places. This is possible because the source key range of any sound is saved when that sound is saved in Audio IFF format. By using Alchemy's resampling function you will be able to tailor any library source sound to almost any sample rate or sample number configuration.

Application 3

Stereo Sampling with Pseudo-Stereo Samplers

Capturing Stereo Sound



Using Alchemy allows you to use three different methods to create stereo sound files. You can cut any section of a true stereo-sampled sound file (a Dyaxis file, for example) and send it to any stereo-playback sampler. This allows you to capture very complex stereo images like ambient stereo reverb and retain them on samplers which can't actually sample in stereo, but can play back in stereo (EPS, Emax, S-900). But to do this you'll need a stereo sampler like the Dyaxis machine to create the source file.

You can also hand-build a stereo file by mixing any waveform anywhere in your stereo image. This lets you do complex panning and image location, but it won't let you create true stereo ambience.

There is a third method which lets you capture true stereo ambience using a mono-sampler. Although it's a bit more complex to accomplish, you can use a mono sampler to sample both channels of a stereo image separately, and then match these channels up in an Alchemy sound file. This practice is particularly useful, for example, if you wish to create a true stereo sound file from a CD sound library, but you only have an Emax, EPS, or S-900. These samplers can easily play stereo sounds back, but using them to capture the stereo sounds requires some planning. Here's how to do it:

- **Choose your stereo source material.**

It's important that you pick a sound with some type of "landmark" (sharp volume peak) which is present in both left and right channels. You will need to use this landmark as a reference point for matching up the left and right waveforms. If the sound has a percussive beginning this will usually be landmark enough, but if the sound has a slow attack you may need to supply your own landmark.

Important Note: If you need to add your own "landmark" you'll probably lose some fidelity on your final sample, because you'll need to create a custom tape version of the sound you're sampling. To do so, hook an alternate sound

source through a mixer with your source playback machine. A drum machine is particularly good for this. Hook both sources up to a high quality tape deck and prepare to record. When you record the copy of your stereo source sound, trigger a short percussive sound just before it begins. Make sure that the percussive sound is mixed evenly in the left and right channels. The tape you've just created is your new source material. It now contains a very visible landmark which you can use to match your channels, and then you can delete it.

- **Hook up the left channel output of your CD player (or tape deck) to the sample input on your sampler.**
- **Make sure that your sampler has enough allocated sampling time to hold the sample you're about to make.**

The sampling rate must be set to a value equal to at least twice the highest frequency you'll be sampling (see About Sound). Use the highest possible sampling rate in order to capture the best sample.

- **Set your sampling threshold so that it will be triggered by your sound's attack (or other "landmark").**

This will assure you that sampling will start as accurately as possible.

- **Arm your sampler and play back the left channel CD source material to sample it.**

Make sure that the quality of the sample you just made is as you desire.

- **Use Alchemy's New command to open a new sound file.**
- **Use the Get Sound command on the Network menu to retrieve the sample you've just made.**

You can review this procedure in detail under Getting a Sound Using Alchemy. After a short transmission period the left channel waveform appears in the untitled window. Now you're ready to sample the right channel.

Constructing a Stereo Sound File

- **Use your sampler to clear the left channel waveform.**

Leave all sample length, sample rate, and sampling threshold settings as they were.

- **Arm your sampler and play back the right channel source material to sample it.**

Make sure that the quality of the sample you just made is as you desire.

- **Use Alchemy's New command to open another new sound file.**
- **Use the Get Sound command on the Network menu to retrieve the sample you've just made.**
- **Choose the Strip command on the Windows menu.**

Your image is now lined up with the left channel in the upper waveform window and the right channel in the lower waveform window.

- **Click on the upper (left channel) waveform window to make it active.**
- **Choose the Soundfile Setup... command on the File menu.**

The Soundfile Setup dialog box appears.

- **Click on the button beside "Stereo" under "Channels."**
- **Click on the OK button.**

You've just split your left channel sound file into a stereo sound file. The waveform in its lower display is just a copy of the upper waveform, so you'll be getting rid of it.

- **Click the mouse in the lower (right) channel of the new stereo waveform window.**

The blinking insertion point appears, and you're ready for editing.

- **Choose the Select All command on the Edit menu.**
- **Choose the Clear command on the Edit menu.**

The lower (right) channel of your stereo sound file is now empty and ready for the true right channel.

- **Click the mouse on the waveform window which contains your true right channel waveform.**
- **Choose the Select All command on the Edit menu.**
- **Choose the Copy command on the Edit menu.**

A copy of the right channel is now on the Clipboard.

- **Click the mouse in the lower (right) channel of the stereo waveform window.**

The blinking insertion point appears, and you're ready for editing.

- **Choose the Select All command on the Edit menu.**
- **Choose the Paste command on the Edit menu.**
- **Click on the OK button of any dialog boxes that may appear and ask about increasing wavesample size.**

Your raw stereo sound file is now complete.

- **Use the Save As... command on the File menu to save your new stereo sound.**

Make sure you save your file in the Audio IFF format, which is the only one that supports stereo sound files. From now on, the saved file will always appear in a stereo waveform window. To match up the channels and finish your task, continue through this section.

Matching Up the Channels of a Stereo Image

- **Choose the Zoom Full In command on the Action menu to see the waveforms as closely as possible.**
- **Use the waveform window scroll bars and arrows to find the waveform of your sound's attack (or "landmark").**

Your stereo channels should be pretty close to being matched up. You want the upper and lower (left and right) waveforms to match up as closely as possible. The two wave shapes will probably be very similar, so it should be easy to see where they naturally match up. You need to make sure that the two waveforms cycle as similarly as possible.

- **Find a cycle in the left channel waveform which definitely matches a cycle in the right channel waveform.**
- **Place the mouse cursor over the channel separator and select a range covering the distance between the cycles' zero crossings.**

The distance between one channel's cycle zero crossing and the other's is the offset which you'll need to correct to match the channels and prevent any phase problems. The channel which has its cycle zero crossing farther to the right is the one which needs to be shortened.

- **Hold down the Option key and use the mouse cursor to slide the selected range to the left so it is before the beginning of the sound**
- **Choose the name of the channel which needs to be shortened from the Windows menu.**

That channel alone is now displayed for editing.

- **Choose the Cut command from the Edit menu.**

You have just corrected your stereo channel offset by cutting out just enough of the lagging waveform to align the cycles you were using as measures. This should assure that the two channels are matched.

- **Choose the Stereo command on the Windows menu.**

Both channels of the sound file are now displayed for editing.

- **Place the mouse cursor over the channel separator and select simultaneously the waveform range in both channels which you want to keep.**

Leave any artificial “landmark” sound and any noise or unwanted waveform range outside of the selected range.

- **Choose the Extract command on the Edit menu.**

The extract dialog appears and asks if it’s OK to get rid of the unwanted waveform sections which are not selected.

- **Click on the OK button.**

The stereo sound file has now been redefined to include only the selected range, and your task is complete.

- **Use the Save command to save your newest version to disk.**

You may also want to send the stereo sound out to your sampler so you can listen to the product. The process for doing this is explained under Sending a Sound in Using Alchemy. When played back on the sampler, the image should be natural, and sound nearly exactly like the source material. If it has a “phased” or “delayed” sound, you may need to tweak it a little more. Once it sounds correct, save it as a final version. From now on it will automatically be assigned to play back in stereo over your chosen key range whenever it is sent to a capable network sampling device.

Application 4

Using Harmonic Analysis/ Resynthesis as a Filtering Tool

Alchemy's FFT analysis and resynthesis functions let you determine the exact harmonic content of a selected waveform range, and then edit that content as you desire. In the Guided Tour section you made some simple alterations to a sound's harmonic spectrum to see how this works. This application gives you an example of how you might use Alchemy's spectrum analysis and resynthesis tools to salvage a sampled sound that is flawed. Perhaps this will give you some ideas about using FFT analysis and resynthesis in your environment.

- **Choose the Open... command on the File menu.**

The open dialog box appears.

- **Insert the Alchemy sound disk in any drive.**

The list of files in the dialog box is updated to show the contents of the Alchemy sound disk.

- **Click on the file "Bad Snare" to highlight it.**
- **Click on the Open button.**

The "Bad Snare" waveform appears in a waveform window on the Alchemy desktop.

- **Click on the Speaker icon to hear what you have.**

The current waveform sounds like a gated snare, but you'll also hear a humming which persists throughout the sound. In its present state the sound is not usable. Your task is to salvage the sound by removing the component of the sound which causes the hum.

- **Double click on the Selection icon to select the entire "Bad Snare" waveform range.**

The "Bad Snare" waveform is now highlighted and ready for editing.

- **Click on the palette's analyze icon.**

This may take a moment because the entire waveform is being analyzed. When the analysis is complete a harmonic spectrum window appears. In it you see a detailed display of the analyzed waveform's harmonic spectrum. Notice the single frequency channel towards the right side of the currently displayed spectrum which pokes up above the others.

- **Click above the highest frequency channel, near the right side of the window.**

A small black square appears at the tip of the channel to indicate that it is selected. Look in the palette's numeric display to see the frequency channel's vital information. Notice that it has a frequency of approximately 175 Hz. It is unlikely that a drum sound with little pure pitch would have such a strong single fundamental. This is probably the pitch which you hear as a hum in your sampler sound.

- Click and hold the mouse on the solid black square at the tip of the selected channel.
- Drag the frequency channel bar downwards until it is approximately even with the channels on either side of it.

Notice how the palette's numeric display shows you the channel's changing amplitude as you do this. By decreasing the height of the selected frequency channel, you will be decreasing the amount of that frequency which will be present in your resynthesized waveform.

- **Click on the "Bad Snare" waveform window to activate it.**

This will let you watch the resynthesis as it happens.

- **Click on the palette's resynthesize icon.**

You're resynthesizing the whole waveform, so this may take a moment. When the waveform window reappears, it contains the new version of the waveform which reflects your harmonic changes.

- **Click on the Speaker icon to hear the resynthesized sound.**

You'll notice that the sound not only looks "cleaner," but it also sounds better. The hum is now gone, so you must have adjusted the correct frequency. If you had adjusted the wrong frequency, you could undo the change with the Undo command, analyze the waveform again, and adjust a different frequency peak.

- **Choose the Undo command on the Edit menu.**

This lets you compare your newly resynthesized waveform to the original. You can see how big the difference really is.

- **Choose the Redo command on the Edit menu.**

Your repaired waveform is now back. You can use the Undo/Redo combination as many times as you like to compare the two. If you wish to keep the salvaged waveform, save it now using the Save command (to save it as "Bad Snare"), or the Save As... command to save it under a different name.



Important Note: When you are dealing with waveforms which consist of more than 32,768 samples, you will not be able to analyze and resynthesize the entire waveform at once. In such cases you can use the EQ... command on the Process menu. This allows you to define very accurate high-pass, low-pass, and notch filtering for waveform ranges of any size. To produce approximately the same end product as you did in the application, you could have used the EQ... command to create a notch filter centered at a frequency of 175 Hz, with a boost/cut of -25.00 dB, and a width of 1 Hz. The EQ... command accomplishes its goal using a different method than the analysis/resynthesis functions. Instead of effecting only the selected frequency to produce a razor-sharp notch, the EQ... command produces a smooth-sided notch similar to hardware equalizers. This means that the amplitudes of frequencies around the notch center are also effected (the amplitude of those closest to the notch center are decreased the most).

Application 5

Working with the SP-1200

Getting a Sound from the SP-1200

The E-mu Systems SP-1200 is a powerful 12-bit sampler designed particularly for creating and playing back percussion samples. Because its function is so specific, its relationship with Alchemy differs from that of standard keyboard-oriented samplers. There are three basic characteristics which you must bear in mind when you're working with the SP-1200: It has a fixed sample playback rate of 26,040 Hz, it has buttons instead of keys for sample assignment, and it allows less outside (Macintosh) control of memory allocation than most samplers. To understand the effects these facts have on the way you'll use your SP-1200, read through the information in the rest of this application. You'll find detailed information about both sending and receiving sounds, as well as directions for resampling waveforms so they'll sound right on the SP-1200.

- **Make sure that your SP-1200 has been added to your network.**

You can find directions for doing this in Adding a Sampler to you Network in Using Alchemy. sampling Once your SP-1200 has been added to your network, it will always appear as a menu choice at the bottom of the Network menu.

- **Choose the New command on the File menu.**

A blank waveform window labeled "Untitled" appears on your screen. This will be the waveform window where your newly retrieved sound will go.

- **Choose the SP-1200 from the bottom of the Network menu.**

A check mark appears in front of it on the menu. You've just told Alchemy you're about to link up with the SP-1200.

- **Choose the Get Sound command on the Network menu.**

After a moment the keyboard dialog appears on your screen. You'll use it to choose which sound you want to retrieve.

- **Click the mouse on the keyboard or use the incremental arrows next to "wave:" to highlight the sound you want.**

Clicking the mouse on the keyboard remote-plays the SP-1200 so you can hear the sound you're about to retrieve. Sounds (waves) are numbered from 0 to 31, and correspond directly to the 32 sounds in your SP-1200. Zero is bank A - sound 1, and 31 is bank D - sound 8. Although the keyboard plays back the sounds, and you can use it to select one for retrieval, you'll find that sound placement is not linear. (The black keys will appear to be numerically out of order.) This is due to the actual MIDI notes assigned by E-mu systems to the internal sounds. For this reason, you may want to use the incremental arrows next to "Wave:" to select your sound. This will step you through sounds 0 through 31 in numerical order. When you have the sound you want highlighted on the keyboard, proceed.

- **Click on the OK button.**

A thermometer-like display appears and shows you the progress of your transfer. When the transfer is complete, the retrieved sound appears in the untitled waveform window where you can edit it and save it.

When you save an SP-1200 sound:

If you save an SP-1200 sound in the Audio IFF format, it will be saved in a file which remembers the sound's source note. This allows you to reopen the file at any time and send it back to your SP-1200 without reassigning it to a pad (key number). If you intend to retrieve an entire SP-1200 full of sounds, it's probably a good idea to store those sounds in their own folder with a particular "SP-1200 drumset name." This will let you open all of the files in the folder at any time and send them to your SP-1200 without reassigning them to new keys.

Resampling an Existing Sound for the SP-1200



- Open the sound file you wish to send to your SP-1200.
- Choose the **Soundfile Setup...** command on the **File** menu.

The Soundfile Setup dialog box appears. Look in the Sample Rate box to determine the rate at which the current sound file was sampled. Any sound you wish to send to your SP-1200 must be sampled at 26,040 Hz. If the sound you have selected to send has a different rate, you'll need to resample. If your sound was sampled at the correct rate, you can skip down to Sending a Sound.

If your sound was sampled at a rate other than 26,040 Hz.:

- Double click on the **Selection** icon to select the entire waveform.

You can also accomplish this by choosing the **Select All** command on the **Edit** menu.

- Choose the **Resample...** command on the **Process** menu.

The Resampling dialog box appears on your screen.

- Click on the button in front of **Sample Rate**.
- Type in the rate **"26040"** in the black box.
- Click on the **OK** button.

A dialog box will appear and ask if it's OK to increase or decrease the number of samples in the sound. A decrease in the number of samples happens when you are decreasing the sample rate. An increase in the number of samples occurs when you are increasing the sample rate.

Important Note: If you are changing the sample rate of a sound to send it to your SP-1200, make sure that you don't accidentally save the sound with its new rate over the original source.



Sending a Sound to the SP-1200

- **Click on the OK button.**

If the sound you're resampling is relatively long, it may take a moment, but when the resampling is done the sound will have its new sampling rate. If you'd like, you can use the Soundfile Setup... command to check this. The sound is now ready to be sent to the SP-1200. If you plan to send this sound to the SP-1200 again, it's probably a good idea to save it now under a new name using the Save As... command on the File menu. This will prevent you from needing to resample the source sound again.

Note: You can actually send a sound which is sampled at a rate other than 26,040 Hz to your SP-1200, but its pitch will be transposed up or down because of the SP-1200's fixed playback rate. This can be used to create some interesting effects.

- **Make sure that the active sound's sample rate is 26040 Hz, and that your SP-1200 has been added to your network.**

Resampling a sound is explained above, and network setup is described near the beginning of the Using Alchemy in the preceeding text

- **Choose SP-1200 at the bottom of the Network menu.**

A check mark appears in front of SP-1200, which tells you that it is the active node on your network.

- **Choose the Send Sound command on the Network menu.**

The keyboard dialog appears.

- **Use the incremental arrows next to "Wave:" to select the destination for your sound.**

Sounds (waves) are numbered from 0 to 31, and correspond directly to the 32 sounds in your SP-1200. Zero is bank A - sound 1, and 31 is bank D - sound 8. Although the keyboard plays back the sounds, and you can use it to select a destina-

tion, you'll find that sound placement is not linear. (The black keys will appear to be numerically out of order.) This is due to the actual MIDI notes assigned by E-mu systems to the internal sounds. For this reason, you may want to use the incremental arrows next to "Wave:" to select your true destination. This will step you through sounds 0 through 31 in numerical order. Then click on the highlighted key to see if any sound is already there. When you have the destination you want highlighted on the keyboard, proceed.

If you wish to send an SP-1200 sound back to its original pad (key) location:

- **Click in the Assign New Voice.button.**

The key assignment of the original sound is automatically highlighted so you can send it back to its origin. This will only work if the sound file has been saved in Audio IFF format (or if you just retrieved it from the SP-1200).

- **Click on the OK button.**

If not enough memory has been allocated on your SP-1200 to place the sound correctly, a dialog box will appear and tell you the length (in seconds) of memory which must be assigned to your chosen destination. *You will need to do this before proceeding.*

If there is enough memory at your selected destination, the transfer will take place automatically, and the transfer thermometer will appear.

To allocate enough memory time to your selected destination:

- **Activate the Sample module on the SP-1200 by pressing the red button.**

The red LED under "Sample" comes on.

- **Press "2" on the SP-1200 keypad to assign a voice.**
- **Select the correct bank and press the destination pad for your sound.**

This sets the SP-1200's Sample module to adjust the data for the correct pad. The SP-1200 asks you whether you want to delete the existing voice.

- **Press the SP-1200's "Yes" button.**

The SP-1200 asks you to select the output channel for the sound you are about to create.

- **Press number you want as your sound's SP-1200 output channel on the SP-1200's numeric keypad.**

This will be the channel over which your sound plays when using the SP-1200's multiple outputs.

- **Press the SP-1200's "Enter" button.**
- **Press "5" on the SP-1200's numeric keypad.**

This allows you to assign the length (in seconds) or memory for your sound. The current length will be listed in the SP-1200's LCD display.

- **Slide the SP-1200 slider number 1 up or down to set the required sound length.**

You can look at the time indicated on your Macintosh screen to make sure of what you need. Set the number in the SP-1200's LCD display to match the Mac number exactly.

- **Press the SP-1200's "Enter" button.**
- **Press "7" on the SP-1200's numeric keypad.**

The sample is now armed.

- **Press "9" on the SP-1200's numeric keypad.**

The SP-1200 tells you that it's sampling, and that the sample is good. The necessary memory is now allocated.

- Press the SP-1200's Sample module button to deactivate sampling and return to normal.
- Click on the OK button in the Mac's dialog box.

The transfer thermometer appears and lets you monitor the transfer. When the thermometer disappears the transfer is complete. A dialog box then appears and gives you exact truncation information for setting the length of the sound.

- Activate the SP-1200's Setup module by pressing the Setup module button.

The red Setup LED lights up to let you know you're in Setup mode.

- Press "19" on the SP-1200's numeric keypad.

You've set your SP-1200 to truncate a sound.

- Press the pad to which you just sent your sound.

This lets the SP-1200 know what sound you wish to truncate.

- Press the SP-1200's "Enter" button.

The sample start (S =), end (E =), and loop (L =) points are now listed in the SP-1200's LCD display. You'll only be adjusting the end point.

- Use the SP-1200's slider 3 and 4 to adjust the end point to match your Mac's dialog box data as closely as possible.

Slider 3 is the rough adjustment and slider 4 is the fine adjustment. If you can't get the exact number indicated by your Mac, choose the highest possible number *which is still below* the suggested point.

- Press the SP-1200's "Enter" button.

The SP-1200 asks you if you want to make the truncation permanent.

- **Press the SP-1200's "Yes" button.**

The truncation has now been made to your sound, and it's ready to use.

- **Deactivate the SP-1200's Setup module by pressing Setup module button.**

The red Setup LED goes out and the SP-1200 returns to its native mode.

- **Click on the OK button in the Mac's truncation dialog box.**

The Alchemy desktop and sound file(s) reappear and you're ready for your next task.

Once you have created a sound file and sent it to your SP-1200, you may wish to save it so you can send it again with the minimum amount of work. To do this, make sure that you use the Save As... command on the File menu to save it in the Audio IFF format. This assures that the file you save will remember the key (pad) assignment of the sound, and stored views you may have used. You also may wish to save any full SP-1200 sets you create to diskettes for quick loading in performance.

Chapter Four

Reference


Introduction

In this section you'll find short explanations of the Alchemy palette and the different windows and commands. You should use this reference section to get quick information about things which are already basically familiar to you. If you are looking for a way to develop this familiarity, try reading through the Guided Tour and Using Alchemy sections of this manual. In those sections you'll find illustrations and step-by-step examples of all program functions.

The Palette

When you start Alchemy, a window containing the palette is always located on the left side of the screen. Learning to use the tools in this window is one key to using Alchemy.

Many Macintosh programs (MacPaint, for example) use palettes. These are usually small windows which contain a number of tool icons and indicators, each representing a process, or the status of a process. By clicking the mouse cursor on an icon, you either prepare to begin a process or you issue a command. In this program, as in MacPaint, using the palette is the fastest and most flexible way to shape what's on the screen.





The Alchemy palette contains seven groups of icons and indicators, and every group focuses on a particular aspect of the program. Although many of these functions can be accomplished from Alchemy's menus, using the palette is generally quicker and easier. Let's look at each of the icon groups and see what they do.

Note: The icons are very specialized, and many represent functions which cannot be performed all of the time. (For example, you can't use the Fade Out icon unless a waveform range is selected.) You can always tell when an icon is not applicable (and therefore disabled), because the icon will be grayed on the palette, and cannot be used.

Mode Icons

At the very top of the palette you see two groups of three icons placed horizontally. These are the mode selection icons, and are the most fundamental of all icons. Any time you are using Alchemy, you'll be in one of these modes.



Selection Mode

Whenever you start up Alchemy, it will be in Selection mode, and the normal mouse arrow cursor will change to an I-Beam whenever it is over any waveform. By holding down the mouse button and dragging the I-beam along a section of waveform, that section is selected and will be outlined in black. You can then click the Speaker icon to hear the selected section. If no range is selected, the entire waveform in the active window is played back. **Shortcut:** Double-clicking on the Selection icon automatically selects the entire active waveform. **Note:** In order for any process to take place (like fading, scaling, or spectrum analysis), you must select a range. The process then takes place only in the selected range. To find out more about this, read through the information about the Process menu.



Draw Mode

When you click on the Waveform Draw mode icon, the active waveform window switches to allow hand drawing of waveforms. In Waveform Draw mode the mouse cursor changes to

a pencil whenever it is over a waveform range. A waveform may then be drawn or altered by dragging the pencil across the window. You must switch back to Selection mode in order to play back a drawn waveform range. Switching to Waveform Draw mode does not change the current resolution of the waveform window, so it's usually a good idea to zoom in before clicking on the Waveform Draw mode icon.



Loop Splice Mode

Clicking on the loop splice icon automatically puts the active waveform window into Loop Splice mode, and zooms in to the splice point between loop end and loop start. Both loop end and loop start can then be adjusted using scroll bars and arrows located at the bottom of the display. It's important to make sure that a loop is turned on before you try to play it back, or the sound won't loop. **Shortcut:** Clicking within the gray scroll bar areas automatically moves the waveform splice point to the next zero crossing.



Amplitude Enveloping Mode

The Amplitude Enveloping icon switches the current waveform window into Amplitude Enveloping mode. In this mode, you can trace and copy existing amplitude envelopes, paste existing waveforms or envelopes over the sounds, or use the mouse to draw or drag a new amplitude envelope for the waveform. Once an envelope has been created or pasted, you can edit the envelope with most of Alchemy's editing functions, and adjust the current sound to fit the new amplitude envelope. For more specific information, see also the Knob/Draw Toggle icon, Amplitude Fit icon, and Amplitude Scale icon.



Frequency Enveloping Mode

The Frequency Enveloping icon switches the current waveform window into Frequency Enveloping mode. In this mode, you can paste existing waveforms or traced envelopes over the sounds for use as frequency modulation curves. It is also possible to draw or drag a new frequency modulation envelope for the waveform. Once a modulation envelope has been created or pasted, you can edit the envelope with most of Alchemy's editing functions, and the instantaneous pitch of the current sound to fit the new modulation envelope. For



more specific information, see also the Knob/Draw Toggle icon, and Frequency Mod icon.

Info Icon (i)

The Info icon performs the same function as the About Alchemy command on the menu. It lists pertinent information about the Alchemy version, and available memory.

Display Icons

Directly below the mode selection icons is a group of six display icons. You will use these icons to change the characteristics of the waveform window display, and to play back the selected waveform range.



Speaker Icon

Clicking inside the Speaker icon will play the currently selected sound or range out of the Mac's internal speaker, or external audio out. To play back a loop, click on the Speaker icon and hold down the mouse button. **Shortcut:** Pressing the space bar automatically clicks the Speaker icon.



Snapshot Icon

Clicking on the Snapshot icon creates an overview display which reflects the current waveform view. It accomplishes this by triggering the Overview function to split the currently selected waveform window into overview (upper), and actual (lower) waveforms. The overview waveform is always located on top, and can be used as a navigation map of the whole sound. A shaded section, or 'gel,' is shown in the overview waveform to indicate what wave section is selected in the current sound file. To update the overview display to reflect the current waveform view, click on the Snapshot icon again.

You can use the overview display to select the wave section you would like to view and edit in the lower waveform display. To do this, simply drag the mouse cursor in the overview (upper) display over the area you would like to view. The area you select will be sized to fit in the lower display,

and a rectangle will appear above the overview waveform showing what section of the waveform you are viewing in the lower display. **Note:** The overview display can not be edited. It is used for view adjustment only.



The Overview Icon

The Overview icon turns on the overview display, which is a map of the whole waveform that can be used as a quick and accurate navigation tool. Clicking on it again while the overview display is showing automatically resets the display to show the full sound. This is especially useful if you have used the Snapshot icon to adjust the overview display to show another waveform view. You can turn off the overview display by choosing the Hide Overview command on the Windows menu.



The Loop Cursors Icon

Click on this icon to turn the sustain loop on and off. When the loop is turned on, two loop cursors appear in the waveform window. These cursors mark the loop start and loop end points, and you can move them by placing the mouse cursor over them, waiting until the cursor turns into a left-right arrow, and then dragging them to new positions. Solid black triangles are displayed at the base of both cursors if axis markers are on, or the Borders preference is chosen using the Preferences... command on the Action menu. The fine positioning of the loop splice point is most easily adjusted in Loop Splice mode. To turn the sustain loop off, click the Loop Cursors icon again.



The Knob/Draw Toggle Icon (enveloping modes active)

This icon replaces the Loop Cursors icon when you are in Amplitude or Frequency Enveloping mode. The icon toggles between the two available envelope editing methods: Knob editing and draw editing. In knob editing, the default setting, you click the mouse cursor on an existing envelope to create 16 break points that can be dragged to define an envelope. There is no limit to the number of break points that can be added. In draw editing, the mouse cursor appears as a pencil, which you can use to hand draw an amplitude or frequency envelope of unlimited intricacy. If you desire, you can combine both envelope editing methods by using knob editing, and then fine tuning with draw editing, or by drawing an

envelope, and “simplifying” it down to a 16 break-point profile. Delete a break point by holding down the option key and clicking on it. Constrain it to the zero crossing by holding down the shift key and dragging. **Note:** When you switch from draw to knob editing, the envelope is automatically reduced to become a 16 point envelope. This cannot be undone, so do it with care.



The Threshold Icon

When you click on the threshold icon, it displays horizontal dotted threshold bars in the waveform window. You can adjust these bars to set an amplitude level for scaling. To set the amplitude level, position the mouse cursor over one of the threshold bars. When the cursor changes to an up-down arrow, click and drag the threshold bars to a new amplitude. Scaling factor and range information are shown in the numeric display. The scaling factor is always based on the highest amplitude present in the selected waveform range. To scale the amplitude of a waveform, select the range to be scaled and click on the Scale icon.



The Axis Markers Icon

Clicking on the Axis Markers icon displays the active waveform window with amplitude and duration rulers. The duration ruler, which is the X axis, may be displayed by decimal sample number, SMPTE frame, or by time in seconds. The amplitude ruler is displayed in percentage of maximum allowable amplitude. Units are selected using the Axis Units pop-up menu on the Windows menu.

Process Icons

Below the display icons is a group of nine process icons which you'll use to accomplish many of your editing tasks. All of the processes represented by these icons can also be executed by using commands on the Process and Edit menus, but you will generally find that using the palette is much quicker and easier. Here are short explanations of each process icon.



The Fade Out Icon

The Fade Out icon fades out the selected waveform range or envelope using the fade slope set by choosing the Edit Options... command on the Edit menu. **Shortcut:** Hold down the command key and click on the Fade Out icon to bring up the edit options



dialog box. You can use this dialog box to select a different fade slope.

The Crossfade Icon

The Crossfade icon automatically mixes a faded out version of the waveform on the Clipboard with a faded in version of the current waveform range, creating a crossfade. The size of the crossfade period is automatically adjusted to fit the length of the shortest waveform. The fade slopes automatically default to the one set using the Edit Options... command on the Edit menu.



The Trace Envelope Icon (enveloping modes active)

The Trace Envelope icon replaces the Crossfade icon when you are in Amplitude or Frequency Enveloping mode. Click on this icon to automatically trace the amplitude envelope of the current waveform. The traced envelope appears as a black line superimposed over the grayed waveform. Once traced, the envelope can be edited using most of Alchemy's waveform editing functions, or copied to the Clipboard for later pasting.



The Fade In Icon

The Fade In icon functions exactly like the Fade Out icon, only the process is reversed. The amplitude of the selected waveform range or envelope is faded from zero to its actual value. Remember, all fades automatically default to the slope set by choosing the Edit Options... command on the Edit menu. **Shortcut:** Hold down the command key and click on the Fade In icon to bring up the edit options dialog box. You can use this dialog box to select a different fade slope.



The Invert Icon

Clicking on the Invert icon turns the selected waveform range or envelope upside down by making all positive sample values negative, and all negative ones positive. Although this doesn't change the way the waveform sounds, it can be very useful for building loops (mirror loops, for example), and for controlling amplitude when mixing waveforms. Clicking on the Invert icon a second time returns the waveform to its original orientation by inverting the sample values once again. **Note:** When you are in Amplitude Enveloping mode, all envelope



values must be positive. In this case, the Invert icon inverts the envelope around the 50% amplitude line, inverting the envelope much like an analog synthesizer would.

The Scale Icon

The Scale icon is generally used in combination with the Threshold Bars icon described under Display Icons. Once you have displayed and adjusted the threshold bars, and selected a range to be scaled, click the Scale icon to execute the amplitude change. The amplitude of all samples in the selected range or envelope will be scaled proportionally to fit under the new threshold. Remember, scaling affects all samples in the selected range or envelope. **Note:** If no threshold bars are displayed, the Scale icon normalizes the amplitude of the waveform in the selected range or envelope, scaling it so that its peaks are set to the 100% level.



The Reverse Icon

When you click on the Reverse icon, the selected waveform range or envelope is reversed, so it will be played backwards when you click the Speaker icon (or press the space bar). Reversing a waveform range is accomplished by trading the first and last sample values, the second and second-from-last values, etc. This is continued until all samples have been swapped. Clicking the Reverse icon again returns the selected range or envelope to its original form by repeating the process.



The Replicate Icon

Using the Replicate icon, you can define a waveform cycle or range, and repeat that cycle or range over a large waveform area. The Replicate icon functions by taking whatever waveform you have cut or copied to the Mac Clipboard and copying it over and over until it fills the selected range.



The Amplitude Scale Icon (enveloping modes active)

When you are in Amplitude Enveloping mode, the Amplitude Scale icon replaces the Replicate icon. Use the Amplitude Scale icon to adjust an existing waveform to a new amplitude envelope. The Amplitude Scale icon can only decrease amplitudes to make the old waveform fit the new amplitude profile. Its effect is less extreme than the Amplitude Fit icon, but better on sounds that contain silence, or large ranges of low amplitude signal.



The Analyze Icon

This icon gives you access to one of Alchemy's most powerful features, Fast Fourier Transform (FFT) spectrum analysis. To use this function, you must first select a waveform range of less than 32,786 samples to analyze. Clicking the Analyze icon will then open up a harmonic spectrum window containing one adjustable frequency channel for each sine wave required to construct the analyzed waveform. Clicking above any bar will display its frequency, amplitude, and phase information in the palette's numeric display. To hear the effect any harmonic edits have on the original waveform, use the Resynthesize icon to rebuild the modified waveform to reflect the new harmonic content.



The Amplitude Fit Icon (enveloping modes active)

When you are in Amplitude Enveloping mode, the Amplitude Fit icon replaces the Analyze icon. Like the Amplitude Scale icon, you can use the Amplitude Fit icon to adjust an existing waveform to a new amplitude envelope. However, the Amplitude Fit icon both increases and decreases amplitudes to make the old waveform fit the new amplitude profile exactly. Its effect is generally more interesting and extreme than the Amplitude Scale icon, but it may cause the addition of noise or distortion on sounds that contain silence, or large ranges of low amplitude signal.



The Resynthesize Icon

You'll find the Resynthesize icon to the right of the Analyze icon. The Analyze and Resynthesize icons always work together. After you have selected a waveform range, analyzed it, and edited its harmonic content in the harmonic spectrum window, you will click the Resynthesize icon to reverse the Fast Fourier Transform and reconstruct the new waveform reflecting your changes.



The Frequency Mod Icon (enveloping modes active)

When you are in Frequency Enveloping mode, the Frequency Mod icon replaces the Resynthesize icon. Use the Frequency Mod icon to modulate the pitch of an existing waveform according to the profile defined by a new modulation envelope. This is the icon you will use to perform all non-linear frequency modulation tasks.

The Waveform View Icons

Below the process icons is a group of three icons over a short horizontal rectangle. These icons are used to change your view in the current waveform or harmonic spectrum window. The Zoom icons and Fit Selection icon are used to change the current view magnification level, while the rectangle indicates the amount of the current waveform which you see in the waveform window. Below are some short explanations of the waveform view icons and indicator.



The Zoom In Icon

Clicking on the Zoom In icon zooms inwards on the left side of the active window to show a smaller section of the waveform, but at a higher resolution. By clicking repeatedly, you can zoom inwards in steps. **Shortcuts:** Holding down the command key and clicking the icon zooms all the way in immediately. Remember also that the overview display offers a host of waveform view possibilities.



The Zoom Out Icon

Clicking on the Zoom Out icon zooms the current waveform display outwards, which squeezes more of the waveform into the active window, but at a lower resolution. By clicking repeatedly, you can zoom outwards in steps. **Shortcut:** Holding down the command key and clicking the icon zooms all the way out immediately.



The Fit Selection Icon

The Fit Selection icon is located to the right of the Zoom Out icon. Clicking on this icon automatically sizes and redraws the currently selected waveform range to fit perfectly in the waveform display for editing.

Cursor Locator Icons

Below the waveform view icons you see the four cursor locator icons, which are grouped into a pair of I-beams and a pair of loop cursors. These icons allow you to move immediately to the beginning or end of the currently selected range or loop, without changing your magnification level. Here are short descriptions of each:



Range Start and Range End Select Icons

The two cursor locate icons on the left are the Range Start and Range End Locator icons. They both look like the mouse I-beam cursor. Clicking on the left I-beam icon automatically centers the view in your waveform window on the beginning of the currently selected range, but it doesn't change the view magnification. The right I-beam icon centers the view on the end of the currently selected range without changing the magnification. **Shortcut:** Holding down the command key when you click either of these icons automatically zooms all the way in around the selected point.



Loop Start and Loop End Select Icons

The two cursor locate icons on the right look like loop cursors (straight black vertical lines). Clicking on the left Loop Cursor icon automatically centers the view in your waveform window on the loop start point, but it doesn't change the view magnification. The right Loop Cursor icon centers the view on the loop end point without changing the magnification. **Shortcut:** Holding down the command key when you click either of these icons automatically zooms all the way in around the selected point. **Note:** Remember that you must turn a loop on in order to display the loop cursors in the active waveform window (see the loop cursor icon under Display Icons).



View Memory Buttons

Below the cursor locate icons is a horizontal rectangle containing the numbers 1 through 8, and a button labeled 'SAVE.' This area of the palette can be used to store eight separate views or ranges for each open waveform window. The view, selected range or insertion point position, and zoom depth are stored. When a memory is recalled the waveform window appears exactly as it was saved. To store your view of any active waveform window, click on the SAVE button and then on the number you prefer. You can recall a stored view for any waveform window by making it the active window, and then clicking on the view number. You can not store overview displays, or loop cursor positions. All stored

X: 676m
Y: 0.81
Sel: 1.910s

views are saved when a waveform window is closed, and will remain as they were saved until they are replaced. **Note:** View memories are only retained by files saved in the Audio IFF format.

Numeric Display Boxes

At the bottom of the palette is a group of five horizontal boxes stacked into a square. The purpose of these boxes depends on the active window.

- When the active window is a waveform window, the top two numeric display boxes show the amplitude and time coordinates of the I-Beam insertion point (or waveform drawing pencil), and the third box shows the size of the selected range in number of samples or time in seconds. You can choose the units for all boxes using the Preferences... command on the Action menu. For more information, see the Axis Makers icon under Display Icons, above. **Note:** When the current waveform window is a stereo window, the top four boxes are used for amplitude and time display.
- When the active window is a harmonic spectrum window, the numeric display boxes show frequency channel number, frequency (in Hz), amplitude, and phase of the last selected frequency channel. Harmonic information can also be displayed linearly or logarithmically by using the Amp dBs and Amp Linear commands on the Windows menu. For more information see The Harmonic Spectrum Window.

The Windows


Alchemy's windows appear as open boxes on the screen, and contain tools and information you can use to analyze, change, and build sounds. The number of windows you can have open at any time depends entirely on the amount of memory your Mac has. On a Macintosh with 1 meg (1024K) of memory, it is quite possible to have over 15 windows open at once, depending on what's in them.

The Waveform Window

In Alchemy, you will see two types of windows. The palette is the vertical window which is always located on the left side of the screen when you start up the program. It contains many icons which are used to manipulate sound in some way. This window can be moved to any position or hidden entirely, but it can't be resized.

All other windows you'll encounter will act exactly like normal Macintosh file windows. You can adjust their size and placement, and move within them using scroll arrows and bars. These windows have one added feature which greatly enhances their usefulness. When the size of a window is changed, its contents are scaled to fit in the new window size. This means that a little window shows you a smaller version of the same picture you see in a giant window.

A waveform window is opened by using the New..., Open..., or Open Special... commands on the File Menu. Once a waveform window is opened, it can be sized to any dimension and all subsequent windows opened will be automatically sized to match. The waveform window displays all of the selected sound's time domain information, going from the sound's beginning on the left, to the sound's end on the right. As with most Macintosh windows, all waveform windows contain scroll bars and arrows. You can use these to move the window around in the displayed sound file.



Stereo Waveform Display


When the Sampler ID preference is chosen, all waveform windows list the name of their current network sampler in the lower left corner of their display. A different network sampler may be chosen by clicking on the up and down arrows next to the current sampler name.

Many waveform windows can be displayed on the screen at once, but only one of these may be active at any time. You can always recognize the active window, because it will have horizontal lines in its 'grabber' bar. When you click the Speaker icon on the tool palette, the waveform in the active window is played back through the Mac's internal speaker or audio out. If a range of that waveform is selected, only that range will play back when you click the Speaker icon.

Performing any process usually requires that a waveform range be selected. Scaling also requires that threshold bars be displayed. For more information about this, see the Using Alchemy section of this manual.

You can create a stereo version of any mono waveform by choosing the Mono to Stereo command on the File menu. You can also change any stereo display back to mono by using the same process, but be careful, because this deletes the right channel.

The stereo waveform display is exactly the same as the mono display, except that two waveforms are shown. The upper waveform is the left channel, and the lower waveform is the right channel. When splitting a mono sound file into stereo, both channels will start out as identical waveforms. Both channels of a stereo sound file may be edited separately, or as a group. To edit them separately, select ranges in a single waveform. To edit the entire stereo image, place the mouse I-beam cursor on the channel separator line which divides the upper waveform from the lower. By clicking and dragging in this area, the same sections are selected in both the upper and lower waveforms.




The stereo waveform display is particularly useful for those with sampling devices designed for stereo sampling and playback, but it is also be a great tool for use with mono samplers that facilitate stereo playback (Ensoniq EPS, E-mu Emax, Akai S900). Stereo playback is accomplished by capturing a true stereo sound from a Dyaxis system, for example, or by sampling both channels of a stereo image separately, matching them up in an Alchemy stereo sound file, and sending the stereo sound file back to the sampler. Alchemy's stereo transfer process automatically places both samples on different layers of the same key range, and pans one to the left and the other to the right. For more information and examples, see the Using Alchemy section of this manual, and the sampler-specific information in the Appendix.

The Enveloping Displays

Whenever you switch from the normal Range Selection mode to Amplitude Enveloping or Frequency Enveloping mode, the standard waveform window is altered to offer you new tools for tracing, copying, pasting, and editing amplitude and frequency envelopes. In both modes, the current waveform becomes gray, indicating that it cannot be edited, and a flat default envelope is shown in black. Bear in mind, however, that these two enveloping displays perform very different tasks.

The Amplitude Enveloping display is a tool for tracing the amplitude envelopes of existing sounds, editing those envelopes, copying and pasting edited or traced envelopes, and adjusting existing sounds to fit new envelopes. The display shows you amplitude on the Y axis and time on the X axis, but you'll notice as you move the mouse cursor over the display that only positive amplitudes are indicated in the numeric display boxes. This is because all amplitude envelope values are positive, and an amplitude envelope never dips below the X axis.


When the Amplitude Enveloping display first appears, a default envelope is shown as a horizontal black line at the top of the display. The envelope is in "full on" position, and is ready to be edited. Once you are looking at the Amplitude Enveloping display, you can create a new envelope for the sound, trace the sound's present envelope, or paste in wave-data or other envelopes. Clicking on the Amplitude Fit or Amplitude Scale icons (or choosing the same commands on the Process menu) then adjusts the sound to fit the new envelope.



The Frequency Enveloping display is a tool for creating pitch modulation curves for existing sounds, and modulating those sounds according to the curve. Like the Amplitude Enveloping display, the Frequency Enveloping display allows you to create and edit envelopes, copy and paste edited or traced envelopes, and modulate existing sounds with new envelopes.

The Harmonic Spectrum Window

The display shows you your pitch modulation range on the Y axis and time on the X axis, and you'll notice as you move the mouse cursor over the display that both positive and negative




modulation values are indicated in the numeric display boxes. This is because any waveform or envelope that is on the Clipboard may be used as the modulation envelope.

When the Frequency Enveloping display first appears, a default envelope is shown as a horizontal black line superimposed over the X axis. The envelope is in “no modulation” position, and is ready to be edited. To modulate the existing waveform, you will first need to choose the modulation range using the Frequency Range pop-up menu on the Process menu. The pitch range you choose on this menu will be indicated on the Y axis when the axis markers are displayed. Choosing a large frequency range will cause extreme pitch modulation, while a small range will cause subtle pitch modulation. Clicking on the Frequency Mod icon (or choosing the same command on the Process menu) then modulates the sound’s pitch according to the displayed envelope.

The Harmonic Spectrum window offers you a method for editing the exact harmonic content of any waveform section. Using the principles of Fast Fourier analysis and resynthesis (outlined in the About Sound section of this manual), the harmonic window offers a very accurate and adjustable

frequency representation of a selected waveform range. Here's how the harmonic spectrum window works:

You must select a waveform range and click on the Analyze icon, or choose the Analyze command on the Process menu. The FFT algorithm is then used to determine the different sine waves which might make up the selected waveform, and their respective amplitudes. When the analysis is complete, a harmonic spectrum window is opened near the bottom of the screen to show the results.




In that display, one vertical bar is shown for each analyzed sine wave frequency. Due to the nature of the FFT algorithm, the number of frequency bands will always be equal to half of the number of overall samples analyzed. Therefore, analysis of larger waveforms will result in a harmonic spectrum containing more frequency channels. At this time, no more than 32,768 samples may be analyzed at one time, which means that no more than 16,384 frequency bands will appear in the harmonic spectrum display. The amplitude of any frequency channel may be adjusted by increasing or decreasing the height (amplitude) of its bar, or deleting it completely. Amplitude may be displayed either logarithmically, or linearly by choosing the Amp dBs or Amp Linear commands on the Windows menu.

When all editing of the harmonic spectrum is complete, a new waveform must always be resynthesized to reflect the new harmonic content. This is accomplished by clicking on the Resynthesize icon, or choosing the Resynthesize command on

The Keyboard Dialog

the Process menu. The FFT algorithm is then reversed to regenerate a new version of the old waveform with the wave shape changes required to reflect the harmonic changes.

After you have configured Alchemy to reflect your sampler or sampling studio, it is possible to send sounds to or get sounds from any sampler in the network. The method for accomplishing this is explained at length in the Using Alchemy section of this manual. Whenever you choose to get or send a new sound, Alchemy's keyboard dialog appears. The keyboard dialog box contains an illustration of the keyboard with grayed keys representing the currently selected keyboard range, a single dark gray key showing the note at which the selected sound was sampled, and a bar indicating the MIDI keyboard zone you're looking at.



Once the keyboard dialog is displayed, you can use the pop-up menus to select the correct sampler, instrument (preset), layer number (for samplers which can trigger more than one sample per key), and wavesample. Choosing a key range to retrieve from is accomplished by clicking the mouse on the correct range. Assigning a new key range and its unity key is accomplished by clicking and dragging across the keyboard. Use the keyboard dialog's octave adjustment scale to slide the keyboard up or down so that it is a window into the correct sampler keyboard range. The bisected circle always marks middle C. A number of the key range assignment tasks can be



automated using the Network menu's Send All command. When all key range selections or assignments have been made, the transfer is started.

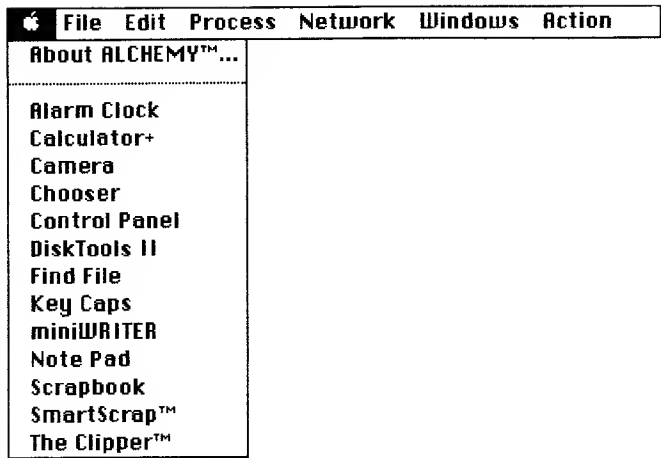
Note: The keyboard dialog also contains a Look button, which can be a very useful tool. If you change the configuration or general state of the selected sampler, you can use the Look button to tell Alchemy that the general information about the sampler's environment should be retrieved again.

In Alchemy, all key range designation, both for sending and retrieving sounds, is handled using the keyboard dialog. Stereo operation is essentially the same as mono operation, except that when you are dealing with stereo sounds the keyboard layer information is handled by the program, and is therefore disabled for the user. Alchemy automatically sends both channels to different layers of the same key range and pans them left and right. For specific information about how to get and send single and multiple sounds, see the Using Alchemy section of this manual.

The Menus

Alchemy, like most Macintosh programs, has a menu bar which remains at the top of the screen at all times. Every word or symbol in the menu bar represents its own menu, each with a number of functions you can perform. To view the different choices on a menu, click the mouse arrow pointer on the menu name in the menu bar. If you hold down the mouse button, the choices are displayed. To execute a command, drag the arrow to the command name while holding down the button. When the desired command is highlighted, letting go of the mouse button executes the command.

The Apple Menu



The Apple menu is represented by the Apple symbol in the upper left of the menu bar. This is a menu which is common to all Macintosh programs, and contains almost exactly the same menu selections in each program. It is generally divided into three sections:

About Alchemy™...

The first selection on almost all Apple menus is an About... command. This usually shows a read-only dialog box which gives background about the program's authors, the release number, etc. This is exactly how the About Alchemy...™ command works. To clear it from the screen, simply click the mouse.

Desk Accessories/MultiFinder

Under the About... command there is a dotted line separating it from your installed desk accessories (DAs). Some standard desk accessories are the calculator, the note pad, and the scrap book. The number of these which appear on your Apple menu depends entirely on how many desk accessories are installed on your system. Most standard desk accessories will function normally within Alchemy. If you are running Alchemy under MultiFinder, you will find a list of the currently loaded programs at the bottom of the Apple menu.

The File Menu

File	Edit	Process	Network	Windows	Action
New					⌘N
Open...					
Open Special...					⌘O
Close					
Save					⌘S
Save As...					
Revert					
Import Resource...					
Export Resource...					
Mono to Stereo					
Soundfile Info...					⌘F
Quit					⌘Q

The File menu is located directly to the right of the Apple menu in the Alchemy menu bar. As the name implies, this is the menu which you will use for most file functions. Here is a list of the different Alchemy file commands, and a short explanation of what each of them does.

New

The New command opens up a blank waveform window and automatically sizes it to match the previously active waveform window. New waveform windows always match the instrument ID and mono/stereo characteristics of the last active waveform window. If no waveform window was active, the blank window is sized to Alchemy's default window size. All waveform windows opened using the New command will be named 'Untitled,' and will be saved that way, unless you use the Save As... command. A blank waveform window can be very useful as a work area for combining and editing existing waveforms. Note: You can display basic file information about any active sound file by choosing the Soundfile Setup command on the File menu.

Open...

The Open... command is used to open existing sound files, and place them in a waveform window so you can edit them. When you choose the Open... command, a dialog box appears on the screen which allows you to choose the folder and file

you wish to open. As always, the new waveform window will automatically be sized to match the previously active waveform window.

Open Special...

The Open Special... command brings up Alchemy's Open Special dialog. This dialog can be used to audition, open, rename, and delete multiple files, and see a listing of pertinent sound file data. By selecting a file and clicking on the listen button, you can hear that sound played back directly from disk. Clicking on the Open button opens the selected file, which appears behind the Open Special dialog. To open an SND resource file, click on the Resource button, and use the pop-up menu to select the resource you want. Use the Open Special... command when you want to open multiple windows without having to choose the Open... command over and over. When you are done opening files, just click on the dialog box's Quit button.

Close

The Close command closes the active window, thereby removing it from the screen. If the active window is a waveform window with unsaved edits, a dialog box will appear and ask you if you want to save the changes before closing. If you don't save your changes before closing, they'll be lost, so be careful. Clicking inside a window's Close box (in the upper left of the title bar) functions in exactly the same manner as choosing this command.

Save

If you choose the Save command, the waveform in the active window will automatically be saved to its original file. Whenever you want to keep a change you've made to any sound, use this command. The Save command should, however, be used with care. If you edit an old sound to create a new sound, the Save command will write the new sound over the old one, deleting the original forever. In cases where you want to keep the old sound *and* the new sound, you should use the Save As... command to give your new sound its own name.

Save As...

The Save As.. command is used to save the waveform in the active window under a name other than the one shown in the window's title bar. When you choose Save As..., a dialog box appears and offers you the chance to name the new file, and select the folder it should go in. The Alchemy Save As...dialog box also lets you save the new file in Audio IFF (16-bit stereo/mono), Sound Designer™ (16-bit mono), Sound Lab™ (8-bit mono), and Dyaxis (16-bit stereo) formats for use with other programs.

If you're not sure which format to use, choose the Audio IFF format, since it guarantees that no 16-bit stereo sound will be translated to a lower-fidelity 8-bit mono sound. It is also the only format which retains the view memories which you've saved within Alchemy. The Sound Designer™ format lets Alchemy owners grab and save in a format compatible with Digidesign's Sound Designer™ program. The Sound Lab™ format is useful for Ensoniq Mirage owners who wish to transfer sounds to or from Blank Software's Sound Lab™ program. The Apple SND Resource format allows programmers to design sounds to be used in Macintosh programs.

Once you have completed the Save As... dialog box, the active waveform window will be saved and labeled with the new name. From now on, you may use the Save command to store your changes to that file automatically.

Revert

The Revert command allows you to go back to the last saved version of the waveform in the active window. If you change a waveform and realize that you liked it better before the change, you can use the Revert command to recall the last version you saved. This allows you to try complex edit combinations, and quickly recall the original waveform if you don't like the result. If you only want undo your last step, you can use the Undo command on the Edit menu.

Import Resource

The Import Resource command brings up an open dialog that allows you to find and open any SND Resource file. SND Resources are the internal Macintosh 8-bit soundfiles used by

the Mac system and various applications (in particular, HyperCard® and HyperCard stacks). To import an SND Resource file, just use the Import Resource dialog to navigate to the application or document that contains that file. When you highlight the application or document, a pop-up menu appears at the bottom of the dialog. Use that pop-up menu to select the SND Resource you wish to import, then click on the Open button. The soundfile then appears in an open waveform window.

Export Resource

The Export Resource command brings up a dialog that allows you to save the current soundfile as an SND Resource file. SND Resources are internal Macintosh 8-bit soundfiles, and are saved directly inside of the applications or documents that will play them. To export an SND Resource file, just use the Export Resource dialog to navigate to the application or document that should contain the file. When you highlight the application or document, a pop-up menu appears at the bottom of the dialog and shows you the SND Resources that are already inside of the selected application or document. To save the current soundfile into the selected application or document, click on the Save Snd button.

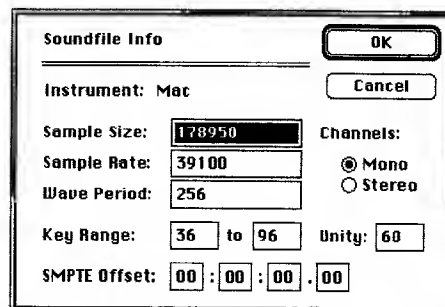
Mono to Stereo/Stereo to Mono

Use the Mono to Stereo and Stereo to Mono commands to convert the current sound file from a one to a two channel sound, or from a two to a one channel sound. When you create a stereo sound from a mono soundfile, the right (lower) channel is automatically filled with a copy of the left (upper) channel. Both channels may be edited separately. When you create a mono sound from a stereo soundfile, the right channel is deleted, so use this command with care.

Soundfile Info...

The Soundfile Info... command opens a dialog box which allows you to view and change all sampling size, rate, period, and key range information for the currently selected sound file. It also lets you set the SMPTE offset, which will be the frame number that is used as the zero time value on the X axis marker.

**The Edit Menu
(waveform
window active)**



A dialog box titled "Soundfile Info" with "OK" and "Cancel" buttons. It contains the following fields:

- Instrument: Mac
- Sample Size: 178950
- Sample Rate: 39100
- Wave Period: 256
- Channels: ☒ Mono, ☐ Stereo
- Key Range: 36 to 96
- Unity: 60
- SMPTE Offset: 00 : 00 : 00 . 00

Quit

The Quit command closes all windows and ends the program, leaving you out on the desktop. For each waveform which you've edited but not saved, a dialog box will appear offering you a last chance to save. If you choose No, the changes you made will be lost forever, so treat each dialog box with care. Generally it's best to save all work before you quit, in which case no dialog boxes will appear when you choose this command, and you will be placed almost immediately on the desktop.

File	Edit	Process	Network	Windows	Action
	Undo		⌘Z		
	Cut		⌘H		
	Copy		⌘C		
	Paste		⌘V		
	Mix		⌘M		
	Insert		⌘I		
	Extract		⌘E		
	Clear				
	Select All		⌘A		
	Select Loop		⌘K		
	Loop Selection		⌘L		
	Auto Zero		⌘U		
	Blending		⌘B		
	Edit Options...				
	Clear Clipboard		⌘R		

You can find the Edit menu directly to the right of the File menu in the Alchemy menu bar. It contains all of the standard Macintosh cut and paste editing functions, as well as a number of new tools which are useful for sound manipulation. Here are the commands you find on the Edit menu, and some short explanations of what they do.

Undo/Redo

The Undo command is an extremely useful editing tool, because it keeps track of your last action and allows you to reverse that action if you don't like the outcome. After you undo something, you can "Redo" it by choosing the command again. This is a great way to compare before and after pictures of any process. Remember, the Undo command only tracks the last action, so use it with caution.

Cut

The Cut command is used to cut a section of waveform and hold that section on the Mac's Clipboard. When a waveform section is cut, the entire waveform area to the right of the cut slides over so that no gap is left. To cut any waveform section, you must first select the desired range. Then choose the Cut command to cut the range to the Clipboard. After a waveform section has been cut, it remains on the Clipboard so that it can be pasted into or mixed with another waveform section. The cut section will remain on the Clipboard until another cut is made, or until the Mac is shut down.

Copy

The Copy command functions similarly to the Cut command, but instead of removing the selected waveform range it leaves the original, and places a copy of it on the Mac's Clipboard. This is useful for many editing techniques, and makes it easy to create new sounds without altering the original. The Copy command is also great for creating echo and reverberation effects by mixing copies of a sound with itself.

Paste

When a waveform range is selected, the Paste command will paste the contents of the Clipboard over the highlighted range. If the Clipboard contents are narrower than the selected range, the paste range will automatically adjust to match the Clip-



board size. If the Clipboard contents are wider than the selected range, the right side of the Clipboard will be truncated (cut off) to fit in the selected range. **Note:** Holding down the command key when you choose the Paste command will paste the Clipboard range starting at the *end* of the selected waveform or harmonic range. This allows you to align a pasted segment with the end of the destination range, as well as the beginning.

If no range is selected, the Paste command will paste the Clipboard contents, regardless of size, over the contents of the active window. The paste will start at the blinking insertion point and continue to the right. The newly pasted range will automatically be highlighted. Remember that a Paste replaces, and therefore deletes, the selected range. If you want to mix the Clipboard range with a new range, use the Mix command on the Edit menu.

Mix

Unlike the Paste command, the Mix command takes the Clipboard contents and adds them to the newly selected waveform range. If the Clipboard contents are narrower than the selected mix range, the mix range will automatically adjust to match the Clipboard size. If the Clipboard contents are wider than the selected mix range, the right side of the Clipboard will be truncated (cut off) to fit in the selected range. **Note:** Holding down the command key when you choose the Mix command will add in the Clipboard range starting at the *end* of the selected waveform or harmonic range. This allows you to align mixes with the end of the destination range, as well as the beginning.

When no mix range is selected, the Mix command will add the Clipboard contents, regardless of size, to the contents of the active window. Just like a paste, the mix will start at the blinking insertion point and continue to the right. The newly mixed range will automatically be highlighted.

Mixing waveforms together combines sounds to produce a single waveform. The percentage of each sound in the resulting mix depends on the gain of the original waveforms.

Insert

The Insert command takes the waveform on the Clipboard and inserts it after the insertion point location in the active waveform window. A space is made in the waveform window so that the Clipboard range will fit exactly, and nothing is deleted. When you select an insert position and choose Insert, a dialog box appears and asks you if it's OK to increase the overall size of the waveform to allow for the inserted section. Clicking OK will proceed with the enlargement. The Clipboard waveform is inserted directly to the right of the insertion point. If a range is selected, the Clipboard waveform is inserted at the left edge of the range.

Extract

The Extract command allows you to select any waveform range and redefine the sound file to consist of only that range. The Extract command executes one-step truncation which can simultaneously delete unwanted waveform data from the beginning and ending of a sound file. Be careful, because this command actually changes the number of samples in the selected sound file.

Clear

The Clear command sets the values of all samples in the selected range to zero, effectively deleting any waveform in that area. The range values are not placed on the Clipboard, so they can not be recalled. Clearing a range does not shorten the sampled sound, it simply changes the selected range to silence. If you want to decrease the number of samples in a sound, you must use either the Extract command to redefine the sound file, or the Resample command to adjust the number of samples.

Select All

The Select All command selects the entire active waveform by selecting all samples in the active waveform window. *Shortcut:* You can accomplish the same thing by double-clicking on the Selection icon.

Select Loop

The Select Loop command selects the defined sustain loop in the active waveform. The loop must be turned on (see Action menu) before the Select Loop command will function.

Loop Selection

The Loop Selection command automatically turns the loop on and places the loop start and loop end cursors at the beginning and end (respectively) of the currently selected waveform range. This is an excellent tool for quickly testing out possible loop ranges.

Auto Zero

When you select the Auto Zero command, a check mark appears in front of it. As long as this check mark is present, all selected ranges will automatically be adjusted to start and end on zero crossings. This is a useful tool for trying out potential loops, and a necessary tool to prevent scaling from creating waveform irregularities. To turn off the automatic zero crossings function, select this command a second time.

Blending

The Blending command is really an option that can be toggled on and off. When a check appears in front of the Blend command, blending is turned on. Blending is an auto-crossfade function that is executed as a result of every Cut, Paste, and Insert command. When you perform one of these edits with the Blend function on, all edit splice points are automatically crossfaded with each other, producing a smooth splice transition. The size of the automatic crossfade range that is used at each splice point is set using the Edit Options... command on the Edit menu. The Blend function is a powerful sound design tool, which (among other things) assures you that your edits will not cause a click, pop, or "brick wall" transition at splice points. Turn off the Blend function by choosing the Blend command again.

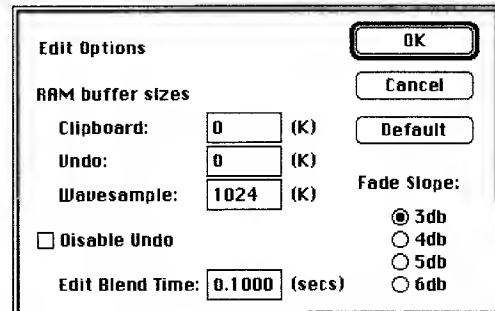


Note: The Blend function accomplishes its task by overlapping the waveforms before and after an edit splice point, according to the size of the Blend Amount (set with the Edit Options command). Since wavedata is being overlapped, the overall duration of the sound will generally be decreased. Keep this in mind when you use the Blend function. (See the Using Alchemy chapter for a diagram explaining the Blend function.)

The Edit Menu (harmonic spectrum window active)

Edit Options

The Edit Options command displays a dialog box which allows you to set the amount of RAM reserved for specific program functions, adjust the Blend Amount crossfade size, and set the fade slopes. Hard disk owners should generally set their Clipboard and Undo buffer sizes to 0.



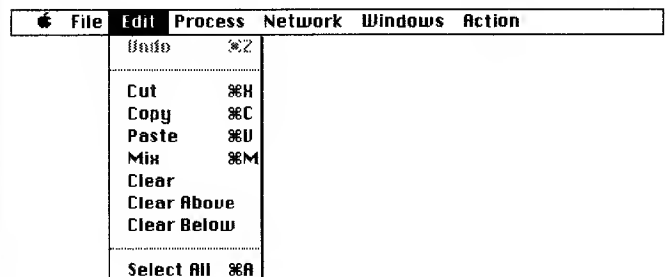
The 'Edit Options' dialog box contains the following controls:

- RAM buffer sizes:**
 - Clipboard: 0 (K)
 - Undo: 0 (K)
 - Wavesample: 1024 (K)
- ☐ Disable Undo
- Edit Blend Time: 0.1000 (secs)
- Fade Slope:**
 - ☒ 3db
 - ☐ 4db
 - ☐ 5db
 - ☐ 6db

Buttons: OK, Cancel, Default.

Clear Clipboard

When you select the Clear Clipboard command, the information that has been cut or copied to the Clipboard is deleted, thereby making more memory available.



File	Edit	Process	Network	Windows	Action
	Undo	⌘Z			
	Cut	⌘H			
	Copy	⌘C			
	Paste	⌘V			
	Mix	⌘M			
	Clear				
	Clear Above				
	Clear Below				
	Select All	⌘A			

You'll find that the Edit menu changes, depending on what type of window is active. The following commands are available when a harmonic spectrum window is the active window.

Undo/Redo

The Undo/Redo command is an extremely useful editing tool, because it keeps track of your last action and allows you to reverse that action if you don't like the outcome. After you undo something, you can 'undo the undo' by choosing the Redo command which takes its place. This is a great way to compare before and after pictures of any process. Remember, the Undo/Redo command only tracks the last action, so use it with caution.

Cut

The Cut command is used to cut a selected section of the harmonic spectrum and hold that section on the Clipboard. When frequency channels of a harmonic spectrum are cut, their dB level is reset to zero, removing those frequencies from the overall harmonic content.

When the harmonic spectrum window is the active window, you can use the mouse pointer to select one or more frequency channels (see the harmonic spectrum window). Once these frequencies are selected, they may be cut using the Cut command. After the frequencies and their levels and phases are cut, they remain on the Clipboard so that it can be pasted over or mixed with other frequency channels. As with a cut waveform, cut frequency channels stay on the Clipboard until another cut is made, or until the Mac is shut down.

Copy

The Copy command functions similarly to the Cut command, but instead of removing the selected harmonic range it leaves the original, and places a copy of it on the Mac's Clipboard. This is particularly useful for certain looping techniques, because it allows you to move frequency channels to another frequency range.

Paste

When a frequency range is selected, the Paste command will paste the contents of the Clipboard over the highlighted range. If the Clipboard contents are narrower than the selected range, the paste range will automatically adjust to match the Clipboard size. If the Clipboard contents are wider than the selected range, the right side of the Clipboard will be truncated

(cut off) to fit in the selected range. **Note:** Holding down the command key when you choose the Paste command will paste the Clipboard range starting at the *end* of the selected waveform or harmonic range. This allows you to align pastes with the end of the destination range, as well as the beginning.

Mix

Unlike the Paste command, the Mix command takes the Clipboard contents and adds them to the newly selected harmonic range. If the Clipboard contents are narrower than the selected mix range, the mix range will automatically adjust to match the Clipboard size. If the Clipboard contents are wider than the selected mix range, the right side of the Clipboard will be truncated (cut off) to fit in the selected range. **Note:** Holding down the command key when you choose the Mix command will add in the Clipboard range starting at the *end* of the selected waveform or harmonic range. This allows you to align mixes with the end of the destination range, as well as the beginning.

Mixing frequency channels together in the harmonic spectrum window has the effect of adding the amplitudes and phases of one frequency range to the amplitudes and phases of another range. This mixes some of the 'harmonic qualities' of the copied (or cut) frequency range with the mix range, and preserves the overall power of the full spectrum.

Clear

The Clear command sets the amplitudes of all frequencies in the selected range to zero, effectively deleting any harmonic information in that area. The range values are not placed on the Clipboard, so they can not be recalled. You should make sure that you want to delete a frequency range permanently before you clear it, because once it's cleared, it's gone.

Clear Above

The Clear Above command functions much like the Clear command, but it automatically clears all frequency bands above the highest selected band. Essentially this is a very accurate low-pass filter which can be used to remove all of the high end from a selected waveform range. This command is only available when a single harmonic channel is selected.

The Edit Menu (enveloping active)

Clear Below

The Clear Below command functions works exactly like the Clear Above command, but it automatically clears all frequency bands below the lowest selected band. Essentially this is a very accurate high-pass filter which can be used to remove all of the low end from a selected waveform range. This command is only available when a single harmonic channel is selected.

File	Edit	Process	Network	Windows	Action
	Undo		⌘Z		
	Copy Envelope		⌘C		
	Paste Envelope		⌘U		
	Clear Envelope				
	Edit Options...				
	Clear Clipboard		⌘R		

The following commands are available on the Edit menu when the active window is an enveloping display.

Undo/Redo

The Undo/Redo command keeps track of your last action and allows you to reverse that action if you don't like the outcome. After you undo something, you can 'undo the undo' by choosing the Redo command which takes its place. This is a great way to compare before and after pictures of any process. Remember, the Undo/Redo command only tracks the last action, so use it with caution.

Copy Envelope

The Copy Envelope command makes a copy of the envelope visible in the current window and places the copy on the Clipboard for later pasting. Copied envelopes may be pasted over other sounds as envelopes, or they may be pasted into a normal waveform window as wavedata, and stored as soundfiles in an envelope library.

Paste Envelope

The Paste Envelope command takes the wavedata on the Clipboard, whether it was an envelope or a waveform, and

The Process Menu (waveform active)

pastes it as an envelope into the current enveloping display. This is an extremely interesting and useful function, because it allows you to use standard sound file waveforms as both amplitude and frequency envelopes.

Clear Envelope

The Clear Envelope command resets the current envelope to the default envelope, which is flat at full amplitude for amplitude enveloping, and flat at zero pitch modulation for frequency enveloping.

Edit Options

The Edit Options command displays a dialog box which allows you to set the amount of RAM reserved for specific program functions, adjust the Blend Amount crossfade size, and set the fade slopes. Hard disk owners should generally set their Clipboard and Undo buffer sizes to 0.

Clear Clipboard

When you select the Clear Clipboard command, the information that has been cut or copied to the Clipboard is deleted, thereby making more memory available.

File	Edit	Process	Network	Windows	Action
		Fade In Fade Out HFade HFade Loop...			
		Invert Scale Reverse Replicate			
		Analyze Resynthesize			
		EQ... Resample... Time Scale... Pitch Shift...			

You'll find Alchemy's Process menu between the Edit and Network menus in the menu bar. In many ways, the Process menu and the palette's process icons represent the power of

this program. Both enable you to perform a number of precise time-domain editing functions on selected waveform sections, and offer you access to FFT frequency analysis/resynthesis tools. It is also important to note that most of Alchemy's processing functions can be used on envelopes, as well as normal sound file waveforms.

There is no difference between executing most of these commands from the palette and selecting them on the Process menu, so you use either method. Over the next few pages you'll find short explanations of each command listed on the Process menu. Remember, these descriptions are listed for reference purposes. To really see these processes in context, you should take a look at the Guided Tour and Using Alchemy sections of this manual. You may also want to read through About Sound to acquire some basic information about the time and frequency domains.

Fade In

The Fade In command performs the same function as the Fade In icon. It fades in the selected waveform range from zero amplitude to 100% of original amplitude over the range area. The fade slope defaults to the one set using the Edit Options... command on the Edit menu.

Fade Out


The Fade out command performs the same function as the Fade Out icon. It fades out the selected waveform range from 100% original amplitude to zero amplitude over the range area. Like all of Alchemy's fade slopes, the fade out fade slope defaults to the one set using the Edit Options... command on the Edit menu.

Xfade

The Xfade command automatically takes the waveform that was most recently cut or copied to the Clipboard and crossfades it with the currently active waveform window. The fade curves default to the settings established by using the Edit Options... on the Edit menu. The crossfade region is automatically adjusted to match the duration of the shortest sound.

Xfade Loop...

The Xfade Loop... command is an automated looping tool that you can use to crossfade your current loop. When you select Xfade Loop..., a dialog box will appear and ask you to set the amount of loop. The fade slopes default to the settings established by using the Edit Options... on the Edit menu. By changing the amount of loop, you adjust the size of the waveform section before loop start which will be mixed with the loop end waveform. A loop amount of 100% will crossfade over the entire loop duration (if possible). For more detailed information about crossfade looping, see About Sound, and Using Alchemy Chapters.



Invert

The Invert command takes the selected waveform range and turns it upside down. This is accomplished by making all positive sample amplitude values negative, and all negative ones positive. Inverting a waveform range will not change its sound in any noticeable way, but it may simplify the creation of certain loops and mixes. Inverting the same range a second time will return it to its original form.

Scale

Use the Scale command to proportionally increase or decrease the amplitudes of all samples in a selected waveform range. By scaling you can adjust the amplitude (volume) of any sampled sound to your liking. To use the scale command, you must select a waveform range and adjust the threshold bars to a new gain threshold. The percentage of scaling to be executed is always indicated in the palette's numeric display next to the word "Scale:." When you select the Scale command (or click on the icon), the selected waveform range will be proportionally adjusted to fit below the new gain threshold.

This can be used to *decrease* the volume of a waveform range (before mixing, for example), as well as *increase* the volume, and is an essential tool for many sound design tasks.

Reverse

The Reverse command takes all samples in the selected waveform range and exchanges the first and last, the second and second-from-last, etc. Essentially, this reverses the order of the samples in the range and causes it to play backwards when you click the Speaker icon. Reversing the same range a second time will return it to its original form.

Replicate

The Replicate command enables you to select a waveform segment to make a repeating copy of that waveform over a large range. This is very useful for constructing a longer waveform from a single wave cycle. The Replicate command functions by copying the waveform which is currently on the Mac Clipboard (the last one you cut or pasted) over and over until it fills the currently selected waveform range.

Analyze


The Analyze command performs Fast Fourier analysis on the selected waveform range. The results of the frequency analysis will appear automatically in the harmonic spectrum window, and can be edited and resynthesized to change the harmonic content of the selected waveform range. Due to the nature of the transform, a harmonic spectrum display will always consist of half the number of frequency bands as samples analyzed. Therefore, larger analyzed waveform ranges will always contain more frequency bands. In order to use the Analyze command, a waveform range of less than 32,768 samples must be selected. For more information, see The Harmonic Spectrum Window.

Resynthesize

Use the Resynthesize command to reconstruct a waveform after you have changed its harmonic content in the harmonic spectrum window. You may resynthesize a waveform as many times as you like until you get the sound you're looking for.


EQ...

The EQ... command is a digital equalization function that allows you to specify particular filtering characteristics and then run the currently selected waveform range through the filter you've created. The EQ... command opens the digital EQ dialog box which allows you to choose to build a notch/peak, high shelf, or low shelf filter with particular frequency center, boost/cut level and frequency width. When the filter is set up, just clicking on the OK button proceeds with the filtering.



Resample...

The Resample... command offers access to a resampling algorithm which is one of Alchemy's most useful functions. You can use it to readjust the sampling rate (or number of samples) in any Alchemy sound file in order to translate sampled sounds from one sampler format to another. When you increase the sampling rate (or number of samples), the algorithm generates new sample values based on analysis of the existing wave. It is not accomplished using a less accurate linear interpolation function. When you decrease the sampling rate (or number of samples), Alchemy uses a skip-sample algorithm to describe the same sampled sound with fewer samples. To resample the currently selected sound file, select the Resample... command, set the new sample rate or sample size in the dialog box, and click on the OK button.



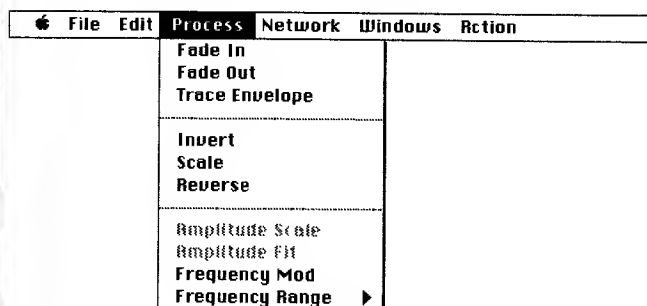
The Process Menu (enveloping active)

Time Scale...

The Time Scale... command opens Alchemy's Time Scale dialog. The Time Scale dialog allows you to adjust the duration of any sampled sound or selected range to a new duration, *without changing its pitch*. The new duration can be set by indicating a new duration, new end time, or a scale factor between the new duration and the old one. Click on the Calculate button to see how one setting affects the others. The time units displayed always correspond to the axis units set with the Axis Units pop-up menu on the Windows menu. You will generally get your best results by using scaling factors which are between 0.7 and 1.3.

Pitch Shift...

The Pitch Shift... command opens Alchemy's Pitch Shift dialog. The Pitch Shift dialog is a keyboard that indicates the current note value and pitch of the sound or range, and allows you to transpose that pitch by clicking on a new key. The amount of the transposition is indicated in the 'Transpose by' box. Click on the box in front of Preserve Duration if you wish the pitch change to leave the sound's duration unchanged. When the pitch shift setting is complete, click on the OK button to execute the shift.



When you are editing an Amplitude or Frequency Enveloping display, the commands on the Process menu change to reflect the available functions. Here are brief explanations of the Process menu commands that are available when for envelope editing.

Fade In

The Fade In command performs the same function as the Fade In icon. It fades in the current envelope from zero amplitude to 100% of original amplitude over its range. The fade slope defaults to the one set using the Edit Options... command on the Edit menu.

Fade Out

The Fade out command performs the same function as the Fade Out icon. It fades out the current envelope from 100% original amplitude to zero amplitude over the range area. Like all of Alchemy's fade slopes, the fade out fade slope defaults to the one set using the Edit Options... command on the Edit menu.

Trace Envelope

The Trace Envelope command automatically extracts the envelope of the current waveform, and displays it on your screen. This is the tool you'll use to isolate the envelope of one sound so it can be superimposed over another sound, or pasted into a normal waveform window and saved as an envelope sound file. Once an envelope has been traced, it can be edited both in knob and draw mode, or processed with the other commands on this menu.

Invert

The Invert command performs two separate functions. If you are editing an amplitude envelope, it takes the current envelope and turns it upside down, while centering it around the 50% amplitude level. This guarantees that no negative amplitude envelope values can be generated. If you are editing a frequency modulation envelope, it takes the current envelope and turns it upside down, exactly as the normal Invert command functions. This is accomplished by making all positive sample amplitude values negative, and all negative ones positive. Inverting the same envelope a second time will return it to its original form.

Scale

Use the Scale command to proportionally increase or decrease the amplitudes of all points in the current envelope. By scaling you can adjust the overall amplitude (volume) of an

envelope to your liking. To use the scale command, you can adjust the threshold bars to a new gain threshold. The percentage of scaling to be executed is always indicated in the palette's numeric display next to the word "Scale:." When you select the Scale command (or click on the icon), the envelope will be proportionally adjusted to fit below the new gain threshold. This can be used to decrease the amplitude of an envelope, as well as *increase* the amplitude. If the threshold bars are not displayed, the scale command will adjust the current envelope so that its peak is at the maximum allowable value.

Reverse

The Reverse command takes all points in the current envelope and exchanges the first and last, the second and second-from-last, etc. Essentially, this reverses the overall amplitude evolution over time. Reversing the same envelope a second time will return it to its original form.

Amplitude Scale

The Amplitude Scale command performs the same function as the Amplitude Scale icon. Use the Amplitude Scale command to adjust an existing waveform to the amplitude envelope indicated as a black line in the Amplitude Enveloping display. The Amplitude Scale command can only decrease amplitudes to make the old waveform fit the new amplitude profile. Its effect is less extreme than the Amplitude Fit icon, but better on sounds that contain silence, or large ranges of low amplitude signal. Remember, the Amplitude Scale command will not increase the overall amplitudes of the current waveform to fit exactly beneath the envelope, but it will still cause the waveform to reflect the current envelope's shape.

Amplitude Fit

The Amplitude Fit command performs the same function as the Amplitude Fit icon. Like the Amplitude Scale command, you can use the Amplitude Fit icon to adjust the displayed waveform to the current amplitude envelope. However, the Amplitude Fit command both increases and decreases amplitudes to make the old waveform fit beneath the new amplitude

envelope exactly. Its effect is generally more interesting and extreme than the Amplitude Scale command, but it may cause the addition of noise or distortion on sounds that contain silence, or large ranges of low amplitude signal.

Frequency Mod

The Frequency Mod command performs the same function as the Frequency Mod icon. Use the Frequency Mod command to modulate the pitch of the displayed waveform according to the profile defined by the current modulation envelope. The degree of modulation is set using the Frequency Range pop-up menu on the Process menu. The Frequency Mod command is the command you will use to perform all non-linear frequency modulation tasks. When you choose this command, the instantaneous pitch of the current waveform is shifted to match the envelope.

Frequency Range

Use the Frequency Range pop-up menu to select the degree of modulation that will be represented on the Y axis of the Frequency Enveloping display. If you choose 2 Octaves, you have set up the Frequency Enveloping display so that the maximum and minimum Y axis values represent pitch shifts of +2 octaves and -2 octaves, respectively. This means that frequency envelopes will cause a great deal of pitch shifting when the Frequency Mod command is chosen. If you choose 1 Semitone, you have set up the Frequency Enveloping display so that the maximum and minimum Y axis values represent pitch shifts of only +1 and -1 semitone, respectively. This means that frequency envelopes will cause a small amount of pitch shifting when the Frequency Mod command is chosen. **Note:** In the Frequency Enveloping display, the Y axis markers are always indicated in semitones.

The Network Menu

File	Edit	Process	Network	Windows	Action
Get Sound ⌘G					
Get Range					
Get All					
Send Sound ⌘D					
Send Range					
Send All					
Instrument ▶					
Mac ⌘1					
Emax ⌘2					
✓EPS ⌘3					

Alchemy's Network menu contains all of the commands you'll use to configure your sampling network to transfer sounds between your Macintosh and your samplers. Since Alchemy works with so many samplers, sample transfer has been generalized. When you first begin using Alchemy you will need to configure the program to understand your network. This is usually accomplished by using the New Instrument... command to add samplers to the Network menu. Once you have defined all of your samplers and their addresses, you'll be ready to select a sampler and get or send sounds. Here are the menu commands you'll use.

Get Sound

The Get Sound command opens the keyboard dialog which allows you to select the sampler, instrument (preset), layer, and key range of the sound you wish to retrieve from the selected network sampler. The keyboard facsimile in the dialog box can be used to remote-play the currently selected sampler for sound choice. Use the pop-up menus to choose the sound you want to get. When you have selected the sound you want, clicking OK will load it to your Mac for editing. If you have retrieved a sound, and then changed the state of your sampler (loaded a new sound, for example), use the Look button to have Alchemy retrieve all pertinent sampler information. Remember that you have to choose the source sampler from the bottom of the Network menu before you can retrieve a sound.

Get Range

The Get Range command functions only with the Emulator III, Roland S-50/S-550, and Ensoniq samplers. It allows you to retrieve a selected range of a waveform which you have already retrieved or sent once with the Get or Send Sound command. This command allows you to skip the keyboard dialog if all you want to do is re-retrieve a particular waveform range. To do so, make sure the correct sampler is selected on the Network menu, select the range, and choose this command.

Get All

The Get All command makes it possible for you to retrieve all of the sounds that make up a particular keyboard configuration. When you choose the Get All command, the keyboard dialog appears. It allows you to select the full keyboard and first sample you want to retrieve from the selected network sampler. Alchemy will begin with the sample you select in the keyboard dialog, and move up the keyboard until all sounds are retrieved (or memory runs out). Each sound will appear in a waveform window as it is retrieved. As usual, the keyboard facsimile in the dialog box can be used to remote-play the currently selected sampler to choose the first sound. Remember that you have to choose the source sampler from the bottom of the Network menu before you can retrieve a sound from it.



Note: If you wish to get or send sounds that have been assigned to key ranges that are outside of your sampler's playable key range (hidden keys), use the octave adjustment scale to adjust the range of keys being viewed.

Send Sound

The Send Sound command functions much like the Get sound command. However, it opens the keyboard dialog which allows you to select the *destination* preset, layer, and key range. The keyboard facsimile in the dialog box can be used to remote-play the currently selected sampler for key range and unity key assignment. If you are assigning a totally new voice to the sampler, check the Assign New Voice box in the keyboard dialog. Then drag to select the new range and unity key. Clicking OK will send the current Alchemy sound file to

the selected sampler and key range. Alchemy files are stored with key range and unity key information, so a sound file will retain its range data. Remember that you have to choose the destination sampler from the bottom of the Network menu before you can send a sound to it.

Send Range

The Send Range command functions only with the Emulator III, Roland S-50/S-550, and Ensoniq samplers. Like the Get Range command, the Send Range command allows you to send a selected range of a waveform which you have already retrieved or sent once with the Get or Send Sound command. This command allows you to skip the keyboard dialog if all you want to do is re-send a particular waveform range. To do so, make sure the correct sampler is selected on the Network menu, select the range, and choose this command.

Send All

The Send All command makes it possible for you to send all of the currently open sound files to the sampler that is selected at the bottom of the Network menu. When you choose the Send All command, a slightly different version of the keyboard dialog appears. It allows you to use pop-up menus to select the full keyboard you want to send all open sounds to. You have three options for defining how you wish to assign key ranges. If you check the box in front of Retain Key Map, the native key ranges of the open sound files will be retained when they are sent to the destination sampler. This is designed for sending multi-sampled sounds from your universal sound library (an entire piano, for example). If you check the box in front of White Keys, the open sound files will be placed on consecutive white keys, beginning at the bottom of the keyboard. This is very useful for trying out drum sounds or sound effects. If you check the box in front of Every "X" Keys, and type in a value for the "X", the open sound files will be placed in consecutive ranges of "X" keys, beginning at the bottom of the keyboard. A send thermometer will appear for each sound that is sent, until all have been sent, or your sampler's memory is full.



Note: You can automatically set all open sound files to the same destination sampler by holding down the option key, and


then selecting the destination sampler's name at the bottom of the Network menu.

Instrument

The Instrument pop-up menu is used to add, adjust, and delete network samplers. It offers these options:

New...

The New Instrument... command opens a dialog box that you can use to add samplers to your Macintosh sampling network. The new instrument dialog box is extremely important, because you will use it to configure Alchemy to communicate with your samplers. It allows you to use pop-up menus to choose sampler type, communication type and channel, Macintosh port, and data transmission speed. The new instrument dialog also allows you to configure Alchemy to control your MIDI patcher. The new instrument dialog must be completed one time for each sampling device in your network. Afterward, the samplers will appear as menu choices at the bottom of the Network menu. To change or delete network sampler information, use the Edit Instrument... or Delete Instrument command.



Edit...

The Edit Instrument... command opens a dialog box that you can use to change the network communications information for existing network samplers. Essentially it's exactly the same as the new instrument dialog box, but you can't change the sampler type, or add a new sampler. It allows you to use

The Windows Menu (waveform active)

pop-up menus to edit communication type and channel, Macintosh port, data transmission speed and MIDI patcher information. All edits change the network as soon as you click the OK button.

Delete

The Delete Instrument is only used to remove a sampler from your network. To execute the command, select the name of the sampler to be deleted at the bottom of the Network menu. Then choose the Delete Instrument command. The sampler is immediately removed from the menu.

Samplers

All network samplers appear as menu choices at the bottom of the Network menu. To communicate with any sampler, you must first select it on this menu. A check mark always appears in front of the currently selected sampler. **Shortcut:** If Alchemy is configured with the Instrument ID preference, each waveform window lists its current sampler in the lower left corner. You can change this display to address another sampler by clicking the mouse cursor on its up or down arrows.

File	Edit	Process	Network	Windows	Action
				Hide Tools	⌘H
				Show Overview	⌘W
				Show Spectrum	
				Tile	⌘T
				Strip	
				Stack	
				Audio Output	▶
				Waveform	▶
				Axis Units	▶
				✓bell door	

The Windows menu contains commands you may use to open and close windows, choose the active window, and select the stereo display you desire. Although many of the commands on the Windows menu can be accomplished by using palette icons or the close boxes on the windows themselves, some

functions such as hiding the overview can only be accomplished by using this menu. Here are some brief descriptions of the commands you find.

Hide/Show Tools

If you choose the Hide Tools command, the palette will be hidden from view. This gives you a larger area of the screen to use for the waveform and harmonic windows, and leaves the program in whatever mode it's in (range select, loop splice, or Waveform Draw mode). If you don't have the Zoom Constraint preference turned on, this command works well with the Zoom box in the upper right corner of every waveform window. If you hide the palette and click on a Zoom box, the selected window will grow to cover the screen, which is great for fine waveform manipulation. This command works like any standard hide/show command. After you hide the palette, the command will change to Show Tools. Choosing it now will make the palette visible again.

Show/Hide Overview

The Show Overview command functions in exactly the same way as the Overview icon, which is described under Display Icons, earlier in this section. Choosing the Show Overview command takes a snapshot of the current waveform and places this snapshot in an overview waveform display at the top of the active waveform window. To turn off the overview display, choose the Hide Overview command, which will appear in place of the Show Overview command. This is the only way to hide the overview.

Show/Hide Spectrum

Choosing the Show Spectrum command provides the same function as the Analyze icon. It opens a harmonic spectrum window which contains a harmonic analysis of the selected range in the active waveform window. The Show Spectrum command will remain dimmed until a waveform range is selected. To close the harmonic spectrum window, choose Hide Spectrum, or click in the window's close box.

Tile

The Tile command automatically takes all of the waveform windows on the screen and evenly covers the screen with them

in a tile-like pattern for easier viewing and editing. This is an excellent command to use immediately after you open multiple files which all overlap. Any tiled waveform can be blown up to cover the screen by clicking on its Zoom box. Clicking in the Zoom box again returns it to its tiled size.

Strip

The Strip command automatically takes all of the waveform windows on the screen and evenly covers the screen with them in horizontal strips, which is particularly good for long sampled sounds. Like Tile, this is a good command to use immediately after you open multiple files which all overlap. Any stripped waveform can be blown up immediately to cover the screen by clicking on its Zoom box. Clicking in the Zoom box again returns it to its stripped size.

Stack

The Stack command automatically takes all of the waveform windows on the screen and places them in descending and overlapping order, just as they would be immediately after you open them. This command is generally useful for people who prefer not to use the zoom boxes, and would rather keep all windows as large as possible.

Waveform

The Waveform pop-up menu lets you decide if you wish to view your waveforms as collections of single sample points, or as connected lines. To choose one, just use the pop-up menu to select the display characteristic you desire.

Axis Units

Use the Axis Units pop-up menu to choose the units you wish to see displayed on the X axis of your waveform windows. You may choose to view the axis in seconds, decimal or hexadecimal samples (for size information), or SMPTE (30 frame). All associated dialogs and tools will use the units you choose here. To choose a unit, just use the pop-up menu to select the axis unit you desire. For related SMPTE information, see the Soundfile Setup... command, on the File menu.

List of open waveform windows

At the bottom of the Windows menu you will always find a

The Windows Menu (harmonic spectrum active)

list of the all open waveform windows, and there will be a check in front of the window which is currently active. To make another waveform window the active window, just choose its title as if it were a command. **Shortcut:** Clicking the mouse anywhere within a window will make it the active window.

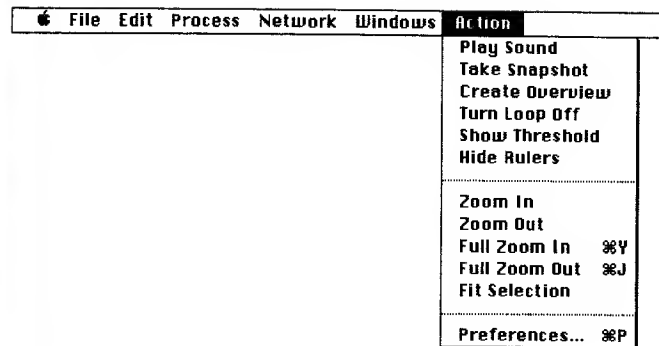
File	Edit	Process	Network	Windows	Action
				Hide Tools	%H
				Show Overview	%W
				Hide Spectrum	
				Tile	%T
				Strip	
				Stack	
				Audio Output	▶
				Harmonic Amp	▶
				✓bell door	

The Windows menu changes according to the type of window which is currently selected. When a harmonic spectrum window is active, the Harmonic Amp pop-up menu is added to the Windows menu.

Harmonic Amp

Use the Harmonic Amp pop-up menu to choose the units you prefer for the Harmonic Spectrum display's amplitude (Y) axis. Choose Amp dBs command to view amplitudes of all frequency channels in the harmonic spectrum window in the non-linear decibel scale. Choose the Amp Linear command, to view the amplitudes of all frequency channels in the harmonic spectrum window using a linear scale, which is often easier for judging relative amplitudes.

The Action menu



The Action menu is the last of the Alchemy menus, and it's located at the far right of the menu bar. The Action menu contains some general display commands, looping commands, and some commands which allow you to customize your program a bit. Here's a list of the Action menu commands, along with some short descriptions of what they do.

Play Sound

The Play Sound command performs the same function as the Speaker icon and the space bar. If a waveform range is selected, it plays the range through the Macintosh's internal speaker or audio out. If no range is selected, it plays back the entire sound in the active waveform window.

Note: If you have a Sound Accelerator™ card installed, the Play Sound command (like the Speaker icon) will play back the selected range or entire sound in 16-bit stereo.

Take Snapshot

If the active waveform window does not have an overview display, the Take Snapshot command takes a picture of the waveform, and places the snapshot in an overview (upper) display. The actual editable waveform is displayed below the snapshot. The overview display is always located on top, and is used as a navigation waveform map. A shaded section, or 'gel,' is shown in the overview waveform to indicate what wave range is currently selected.



Once the active waveform window has an overview display, you can update the overview waveform to reflect the current waveform in the lower waveform view by choosing the Take Snapshot command again. The overview display can be updated this way as many times as you like.

You can also use the overview waveform to select the wave section you would like to view in the lower waveform display. To do this, simply use the mouse cursor to select a range in the overview (upper) display. The area you select will be sized to fit in the lower display, and a rectangle will appear above the overview waveform showing what section of the waveform you are viewing in the lower display.

Create/Hide Overview

The Create Overview display always places an upper overview display in the selected waveform window. The overview shows a map of the entire current waveform. When you use the mouse to select a waveform range in the overview display, that range is automatically sized to fit in the lower waveform display where it can be edited. The Hide Overview command removes the overview display from the current waveform window. This is the only way to hide an overview display.

Turn Loop On/Off

The Turn Loop On/Off command turns the loop on if it's off, and off if it's on. When the loop is turned on, two loop cursors appear in the waveform window. These cursors mark the loop start and loop end points, and you can move them by dragging them to new positions. The fine positioning of the loop splice point is most easily adjusted in Loop Splice mode.

Show/Hide Threshold

The Show Threshold command displays the horizontal dotted threshold bars in the waveform window. You can adjust these bars to set an amplitude level for scaling. To set the amplitude level, position the mouse cursor over one of the threshold bars. When the cursor changes to a bi-directional arrow, click and drag the threshold bars to a new amplitude. The scaling amount is always shown in the palette's numeric display. To execute the amplitude change, make sure that you've selected

a waveform range, and choose the Scale command. (You can do this either from the Process menu, or from the palette's process icons).

Show/Hide Rulers

When you choose the Show Rulers command, the active waveform window will be displayed with amplitude and duration rulers. The duration ruler, which is the X axis, may be displayed by decimal sample number or by time in seconds. The amplitude ruler is displayed in percentage of maximum allowable amplitude. To change the display units use the Preferences... command on the Action menu. To turn the rulers off, choose the Hide Rulers command or click on the Axis Markers icon a second time.

Zoom In

Choosing the Zoom In command is the same thing as clicking on the Zoom In icon. It zooms the current waveform or harmonic view inwards on the left side of the active window to show a smaller section of the waveform or spectrum, but at a higher resolution. By selecting this command repeatedly, you can zoom inwards in steps. **Shortcut:** Holding down the command key and clicking on the Zoom In icon zooms all the way in immediately.

Zoom Out

The Zoom Out icon is located to the right of the zoom in. Choosing the zoom out command zooms the current waveform display outward, which squeezes more of the waveform into the active window, but at a lower resolution. By choosing the command repeatedly, you can zoom outwards in steps. **Shortcut:** Holding down the command key and clicking on the Zoom Out icon zooms all the way out immediately.

Full Zoom In

The Full Zoom In command automatically zooms the current waveform display all the way in at the left edge of the window. You can accomplish the same thing by holding down the command key and clicking on the Zoom In icon.



Full Zoom Out

The Full Zoom Out command automatically zooms the current waveform display all the way out to show the entire active waveform. You can accomplish the same thing by holding down the command key and clicking on the Zoom Out icon.

Fit Selection

The Fit Selection command automatically sizes and redraws the currently selected waveform range to fit perfectly in the waveform display for editing.

Audio Output

Use the Audio Output pop-up menu to set your Mac's audio playback status. Your Macintosh's sample playback rate can be set to the default rate (22,257 Hz), or the wavesample's actual (variable) sample rate. The Mac's fixed sample playback rate is 22,257 Hz, so it has to use algorithms to emulate different rates. The true pitch of any sample is only reflected when it is played back at its original sampling rate. For this reason, the wavesample's actual rate is generally preferable. The Sound Accelerator option is only available to those Alchemy owners who have installed a Sound Accelerator™ card, in order to play back sounds in true 16-bit stereo.

Note: This command only adjusts how the sound is played back through the Mac, and does not effect the sample data in any way.

Preferences...

The Preferences... command opens a dialog box which lets you choose these basic options concerning how Alchemy will look and function. The following options can be set:

Preferences

OK Cancel

Startup: Environment:

☐ New ☒ Zero Crossing ☐ Borders

☒ Open ☒ Channel Separator ☒ Zoom Constraint

☐ None ☒ Instrument ID ☒ Auto Resample

Numeric Values: ☒ Decimal ☐ Hexadecimal

Zero Crossings: Displays a black line representing the zero line in waveform windows.

Channel Separator: Displays a black separator line between right and left channels in stereo waveform displays.

Instrument ID: Shows an adjustable display in the lower left corner of all waveform windows which lists their current network sampler source/destination.

Borders: Leaves open space above and below waveforms to show tabs at loop cursor bottoms and to facilitate easy grabbing of threshold bars.

Zoom Constraint: Prevents the Tile and Strip commands from placing waveform windows underneath the palette.

Auto Resample: Automatically resamples sounds as you send them, so that their sample rates match a sample rate that is actually available on the destination sampler. If necessary, you will be prompted to decide whether the sound should be upsampled (rate increased) or downsampled (rate decreased).

The preferences dialog also lets you choose to sample numbers in decimal or hexadecimal format. See the Axis Units command on the Windows menu for more information. To set all preferences, select the desired settings and click OK. From then on, Alchemy will remember your preference information every time you start up.

Appendix

MIDI and SCSI: Communications Standards

Once you have familiarized yourself with some of the basic concepts behind the creation and digitization of sound, it's a wise idea to take a look at how the machines responsible for this task communicate with each other.

Communication wouldn't really be a concern if sampling, editing, looping, processing, and playback were all handled by a single machine, but in our less-than-ideal world this is not the case. Although most instruments which are used to play back samples can accomplish many of these tasks, they don't always do it best or most accurately.

Sampling keyboards and drum machines may be the most efficient playback devices, but dedicated computers are best for editing the data. Hardware instruments are compact and sturdy, and are usually constructed to recreate sampled sounds polyphonically and with high fidelity. Computers are much more flexible, usually have expandable memory, and are built to manipulate data as quickly and flexibly as possible. Computers have been used in musical applications for a long time, but thanks to the sinking cost of personal computers and the acceptance of a communications standard called MIDI (Musical Instrument Digital Interface), computers are currently being used more than ever for musical tasks.

Although MIDI was developed after 1980, you would now be hard pressed to find an electronic instrument that doesn't have it. It was originally designed to provide a way for electronic instruments to talk to other electronic instruments. By plugging the instruments' MIDI ports together with a MIDI cable, one instrument could be used to choose and remote-play sounds on the other. This allowed many instruments to be controlled from a single MIDI controller, and truly simplified life for many musicians.

As time passed, the role of MIDI expanded. Instruments were constructed to dump all types of information over MIDI to a computer with a MIDI interface. There, it was much easier to see what you were doing. The computer could be used to store a seemingly endless number of synthesizer patches, so its role as a sound librarian grew. Unfortunately, the designers of MIDI had no idea how far MIDI would go. When inexpensive

memory and sampling capabilities became available, the computer's use as visual sample editor and real-time MIDI sequencer began to make it indispensable to many musicians.

Although sequencing often seemed to put a strain on MIDI, sampling is where the MIDI standard was pressed past its limit. MIDI is a 10-bit serial communications standard. If you think back to the discussion of bit formats, you'll understand what this means. In an 8-bit communication, every "word" is eight bits of information long. Serial communication transmits this word one bit at a time. After the tenth bit is sent, the first bit of the next word is sent, and on it goes. The receiving device knows that the words will be coming in serial form, and breaks them into 10-bit sections to get the original words back. MIDI was designed to send 31,250 bits per second. That means that it can only send 3,125 10-bit words per second (or 1562.5 16-bit words per second).

When you remember that an average sampling rate might be 50,000 samples per second (50 kHz), with each sample requiring a word of information, you can see the problem. To sample using 8 bits directly over MIDI, you would need to send at least 400,000 bits of serial information per second (50,000 samples multiplied by 8 bits). There is no way to accommodate this without radically changing MIDI, and it seems too late for that.

This doesn't mean that MIDI isn't useful. For sequencing, librarian work, and remote keyboard control it is wonderful. When a sound has been sampled, MIDI, though not fast, is great for dumping that sampled sound to a computer. What MIDI can't do is pass sampled information as it's happening. This is a major limitation. If it could, then almost any home computer could be used as a high-quality direct-to-disk digital recorder. You would accomplish this by storing the samples directly on a computer's hard disk as you took them. Unfortunately, you can't do this over MIDI.

SCSI (Small Computer System Interface) gets around MIDI's problems in two ways. First, it is a parallel communications standard, which means that it can send a number of bits simultaneously. SCSI actually sends eight bits at a time,

which is one word on an 8-bit machine, or 1/2 word on a 16-bit machine. You can visualize this by picturing eight wires, each carrying one bit of the word. Every 8-bit word goes in whole and comes out whole.

Not only does SCSI send eight bits at a time, but it does so even faster than MIDI sends a single bit. Where MIDI sent 3,125 10-bit words per second, SCSI can send between 250,000 and 1,500,000 8-bit words per second. This is well above the minimum speed required to send quality samples in real time (as they happen), both for 8-bit, and larger-word machines.

So theoretically, SCSI picks up where MIDI leaves off, and some newer sampling instruments are being built with both MIDI and SCSI ports. As you would expect, SCSI is not the answer to all problems, but it does offer the speed and compatibility required for the direct-to-disk sampling, which is the heart of the tapeless studio. It may never replace MIDI, but with time it will certainly take its place in the sampling musician's studio.

Ensoniq EPS Specifics

Compatible Software Version:
All.

Communications Types:
MIDI.

How is Stereo Achieved?

When a stereo sound file is created and sent to an Ensoniq EPS, Alchemy automatically places the left and right channel waveforms on two separate layers and assigns them to the same key range. One layer is then panned to the left and the other is panned to the right. Alchemy always places the left channel on the current layer and then creates another layer (which is one number higher) for the right channel. To place the stereo image on layers 7 and 8, for example, you simply need to create layer 7. Alchemy automatically places the left channel waveform there, and then creates layer 8 for the right channel.

Other Information:

- For Alchemy to function with your EPS, you will need to set the System Exclusive parameter to "ON." It can be found on the MIDI page of the operating system. Then save global parameters. This assures you that your EPS will always communicate with Alchemy.
- As with many samplers, exact sample rates on the EPS are not always evident. If you plan to resample library sounds and send them to an EPS, you will need to know the exact EPS Sample rate before you can accomplish the task. One way to figure out any exact sample rate is to retrieve a sound from the EPS that is sampled at the approximate rate you desire. Then choose the Soundfile Setup... command on the File menu to find out the exact sample rate of that sound.
- Along with the Mirage, the Ensoniq EPS allows any selected waveform range to be sent and retrieved. This is accomplished using the Send Range and Get Range commands on the Network menu.

E-mu Emax Specifics

Compatible Software Version:
3.2 or above.

Communications Types:
MIDI or RS-422.

How is Stereo Achieved?

When a stereo sound file is created and sent to an E-mu Emax, Alchemy automatically places the left channel waveform on the primary layer and the right channel waveform on the secondary layer, and assigns them to the same key range. Primary is then panned to the left and secondary is panned to the right.

Other Information:

- The RS-422 communications standard is a high speed communications type which greatly accelerates the transfer of wave data.
- As with many samplers, exact sample rates on the Emax are not always evident. If you plan to resample library sounds and send them to an Emax, you will need to know the exact Emax sample rate before you can accomplish the task. (For example, an Emax rate of "28 kHz" is actually 27778 Hz.) One way to figure out any exact sample rate is to retrieve a sound from the Emax that is sampled at the approximate rate you desire. Then choose the Soundfile Setup... command on the file menu to find out the exact sample rate of that sound.

Akai S900 Specifics

Compatible Software Version:

All.

Communications Types:

MIDI.

How is Stereo Achieved?

When a stereo sound file is created and sent to an Akai S-900, Alchemy automatically creates two key groups with the same defined key ranges. The left channel waveform is sent to one key group, and the right channel waveform to the other. Then both channels are panned to their correct sides in the stereo image.

Other Information:

- On the Akai S-900, sample rates are adjusted by bandwidth, which can make it difficult to figure out the exact sample rate in Hertz. If you plan to resample library sounds and send them to an S-900, you will need to know the selected Akai sample rate in Hertz before you can accomplish the task. One way to figure out any sample rate is to retrieve a sound from the Akai that is sample at the bandwidth you desire. Then choose the Soundfile Setup... command on the File menu to find out the exact sample rate in Hertz.

E-mu SP1200 Specifics

Compatible Software Version:

June 11 21:09 or later.

Communications Types:

MIDI.

How is Stereo Achieved?

Not possible.

Other Information:

- The Disk Module and Sample Module must not be active (their LEDs should not be on) during transfer. This will prevent the transfer from taking place.
- When the SP-1200 is the selected network node, the Keyboard window will display pad numbers (in the A1 to D8 format) as the Sample Name. This makes it easier to keep track of the actual SP-1200 pad that you're addressing.
- The SP-1200 does not allow allocation of sample memory by an external source, so all sample and truncation times must be set using the SP-1200 itself.
- The SP-1200 has a fixed sample rate of 26,040 Hz., so all sounds must be resampled to this rate in order to maintain original pitch. For more information, see Working with the SP-1200 in the Applications section of Using Alchemy.

Ensoniq Mirage Specifics

Compatible Software Version:
Any MASOS.

Communications Types:
MIDI.

How is Stereo Achieved?
Not possible.

Other Information:

- The Mirage does not easily allow allocation of sample memory by an external source, so all sample times must be set using the Mirage itself.
- The Mirage's loop start point resolution is 256 samples, so Alchemy will often be forced to adjust your selected loop start point to the nearest allowable sample position.
- When the Mirage is the selected network nod, the Keyboard window will display the Mirage Program number (1-4) as the Instrument. This allows you to assign different key ranges to the four different Mirage programs.

CASIO FZ-1/ FZ-10 Specifics

Compatible Software Version:

All.

Communications Types:

MIDI.

How is Stereo Achieved?

Four voice polyphonic stereo playback is achieved on the FZ-1 and FZ-10 by routing the left channel through outputs 1-4, and the right channel through outputs 5-8.

General Information:

- To send and receive MIDI wave data with the FZ-1 and FZ-10, you must follow this procedure:

- 1) Press the MODIFY key.
- 2) From the MAIN MENU, use the cursor keys to point at DATA DUMP. Press ENTER.
- 3) From the DATA DUMP screen, use the cursor keys to point at SELECT DEVICE. Press ENTER.
- 4) Press the YES key twice to select MIDI as the DUMP DEV.
- 5) Press the down cursor arrow to point at REMOTE MODE.
- 6) Press the CALL/SET MENU button to allow quick access back to the SELECT DEVICE screen.

You are now ready to retrieve from or send to the FZ-1 or FZ-10. In order to transmit wave data to or from the FAZ samplers, you must always be on the SELECT DEVICE page with the arrow cursor pointing at REMOTE MODE.

- When using the Alchemy keyboard dialog you must press the PLAY button before you can change from one FZ bank (program) to another. The PLAY screen is the only one which allows program changes. You can return to the SELECT DEVICE screen by pressing the CALL/SET MENU button.

- Due to the design of the FZ-1 and FZ-10, loop point changes are not sent remotely to the samplers. The correct loop-points are only sent when the entire sound is sent.

- If you use the FZ-1 or FZ-10 operating systems to delete sounds while you are using Alchemy, you run the risk of upsetting the keyboard placement of newly sent sounds. If you must use the sampler's operating system to delete a voice during an Alchemy session, make sure to **SAVE** your changes and then **LOAD** them back into the sampler before sending a new sound.
- If you retrieve a mono FZ-1 or FZ-10 sound, split it into stereo, and then send it back to replace the original, your stereo image will not be correctly routed to the outputs. When creating a stereo sound from a mono sound, it is better to assign the new stereo sound as if both channels are new voices. This will assure you that the stereo image will be routed correctly.
- The nature of the FZ-1 and FZ-10 is such that the top key of any key range cannot be higher than three octaves above the unity key.

Roland S-550 Specifics

Compatible Software Version:
All.

Communications Types:
MIDI.

How is Stereo Achieved?

Two separate sampled sounds are sent to the first and second tones of a single patch. They are automatically panned so that the left channel plays through output 1 and the right channel plays through output 2.

General Information:

- To send and receive MIDI wave data with the S-550, you must make sure that the System Exclusive parameter is set to ON.
- When you replace an old sampled sound with a new, longer one, the S-550 does not reset the length to hold the entire new sound. A better method is to assign any new sounds, or longer versions of an old sound using the Assign New Voice option in Alchemy's keyboard dialog. Then delete the old sound, if necessary.
- In the S-550, samples sounds (tones) are often shared by many different keyboard configurations. For this reason you should delete sounds with care, or you may find that your action is more far-reaching than you desire.

Roland S-50 Specifics

Compatible Software Version:
All.

Communications Types:
MIDI.

How is Stereo Achieved?

Stereo is not possible on the S-50, but the Alchemy keyboard dialog and stereo display can be used to edit keyboard layer information.

General Information:

- To send and receive wave data with the S-50, you must make sure that the System Exclusive parameter is set to ON.
- When you replace an old sampled sound with a new, longer one, the S-50 does not reset the length to hold the entire new sound. A better method is to assign any new sounds, or longer versions of old sounds, using the Assign New Voice option in Alchemy's keyboard dialog. Then delete the old sound, if necessary.
- In the S-50, sampled sounds (tones) are often shared by many different keyboard configurations. For this reason you should delete sounds with care, or you may find that your action is more far-reaching than you desire.

SAMPLE DUMP: Sequential Prophet 2000 and 2002 Specifics

Compatible Software Version:
All.

Communications Types:
MIDI.

How is Stereo Achieved?

Stereo can only be achieved on SDS devices by sending both channels separately.

General Information:

- The MIDI Sample Dump Standard (SDS) is a stripped-down communications standard which allows the transfer of wave data from machine to machine. The SDS format contains information about sample rate, loop points and raw wave data. It does not contain any information about stereo configuration, key range, or unity key. All SDS wave sample you retrieve will require the assignment of a new key range and unity key, and their sustain loops will need to be turned on using the Loop Cursors icon on Alchemy's palette. When sending a wave sample using the SDS, you will always choose the placement of that sample sound in the destination machines's memory by using a numeric dialog box. All other parameters must be set by adjusting the destination device.
- When the 2000 or 2002 send their last block of wave data, you will receive an **ERROR ON LAST BLOCK** message. You may disregard this, as it has no effect on the sound.
- Neither the 2000 or the 2002 will automatically allocate memory for a new wave sample. For this reason, you should allocate memory manually on the samplers before sending a sound. If you do not do this, longer sounds will be truncated and shorter sounds will retain remnants of the old sound at the end.

SAMPLE DUMP: Yamaha TX 16W Specifics

Compatible Software Version:
All.

Communications Types:
MIDI.

How is Stereo Achieved?

Stereo can only be achieved on SDS devices by sending both channels separately.

General Information:

- The MIDI Sample Dump Standard (SDS) is a stripped-down communications standard which allows the transfer of wave data from machine to machine. The SDS format contains information about sample rate, loop points and raw wave data. It does not contain any information about stereo configuration, key range, or unity key. All SDS wave samples you retrieve will require the assignment of a new key range and unity key, and their sustain loops will need to be turned on using the Loop Cursors icon on Alchemy's palette. When sending a wave sample using the SDS, you will always choose the placement of that sampled sound in the destination machines' memory by using the numeric dialog box. All other parameters must be set by adjusting the destination device.
- The TX 16W only accepts wave data up to loop end, and then aborts. This is a function of the TX16W's design, which requires that the loop end is always the same as the sample end.
- The TX16W will not allow you to use the Sample Dump Standard to replace old samples with new ones. All new wave samples are automatically assigned to the next available wave sample slot.

SAMPLE DUMP: Oberheim DPX-1 Specifics

Compatible Software Version:
All.

Communications Types:
MIDI.

How is Stereo Achieved?

Stereo can only be achieved on SDS devices by sending both channels separately.

General Information:

- The MIDI Sample Dump Standard (SDS) is a stripped-down communications standard which allows the transfer of wave data from machine to machine. The SDS format contains information about sample rate, loop points and raw wave data. It does not contain any information about stereo configuration, key range, or unity key. All SDS wave samples you retrieve will require the assignment of a new key range and unity key, and their sustain loops will need to be turned on using the Loop Cursors icon on Alchemy's palette. When sending a wave sample using the SDS, you will always choose the placement of that sampled sound in the destination machines' memory by using a numeric dialog box. All other parameters must be set by adjusting the destination device.
- Sending a sound to the DPX-1 using the Sample Dump Standard always erases all other sounds in memory. For this reason, you will need to send each sound separately to the DPX-1, and then save it to disk.
- Retrieving a Mirage sound from the DPX-1 using the Sample Dump Standard often results in an incorrect sample rate. This is a function of the Mirage's architecture, which does not include the sample rate as part of the transmitted waved data. To remedy this, you can adjust the sample rate in Alchemy by using the Soundfile Setup... command on the File menu.

Ensoniq EPS/SCSI Specifics



Compatible Software Version:

Version 2.2 or higher (SCSI card and SCSI ROMs required).

Communications Types:

SCSI (MIDI connection is also required).

Boot-up Procedure:

- Turn on all external drives, including CD-ROMs.
- Boot the EPS with software version 2.2 or higher. Make sure that the message "SCSI Installed" appears on the EPS display.
- Boot your Macintosh.
- If your EPS did not display the message "SCSI Installed," then reboot your EPS.
- You should now be ready to use Alchemy .

Note: If the Macintosh or the EPS seem to lock-up repeatedly during the boot procedure, you may have duplicate SCSI ID's (see the SCSI ID Information Table). If nothing seems to function correctly, try booting all devices first, and then hooking up your EPS to Mac SCSI cable. Use this only as a last resort.

EPS SCSI Instrument Definition

Use Alchemy's New Instrument command on the Network menu to add a SCSI EPS just as you would add any new network instrument. There is, however, a major consideration when adding a SCSI EPS: The EPS has a fixed SCSI ID# of 3. This can not be changed, so make sure that any SCSI EPS is set to this SCSI ID number.

Sync Mode

When running Alchemy with the EPS, a special mode option appears in the Network menu called "Sync Mode." In sync mode, when you execute an Alchemy editing or processing command, the edited part of the sound will automatically be transferred to the EPS. The result of this is that the sound in the Macintosh and the sound in the EPS stay in sync (are the same), without having to choose the Send Sound command. This type of operation is not practical without the high-speed benefit of SCSI data transfers.

Sync mode can only remain active when the sound you're working on in Alchemy still resides in the memory of the EPS. For example, if you load a number of sounds into the EPS and transfer one over to Alchemy for editing, a link between the Macintosh and EPS is established, and sync mode will be operative. If you then erase that bank or load a new sound into the EPS) the link to the sound in Alchemy will be severed, and sync mode will be inoperative.

If you wish to isolate the editing performed on the Macintosh from the EPS, you can turn sync mode off by choosing it again on the Network menu.

A Word About SCSI ID's

Before the EPS and Alchemy can boot-up and operate properly, you must be certain that every device in your SCSI chain has a unique SCSI ID#. Available SCSI addresses range from 0 to 7, but **the EPS is set to ID#3 at the factory, and cannot be altered.** Most SCSI hard disks and peripheral manufacturers allow you to select SCSI ID numbers manually, to avoid potential conflicts. For a list of some standard SCSI ID numbers, see the SCSI IDInformation Table.

E-mu EIII Specifics



Compatible Software Version:
Version 2.0 or later.

Communications Types
SCSI.

Boot -up Procedure:

- Turn on any external EIII drives such as the Emu HD 300.
- Boot the EIII with software 2.0 software as indicated on the EIII system disk label.
- Turn on any external Macintosh hard drives and devices (i.e., CD-ROM).
- Boot your Macintosh.
- You should now be able to use Alchemy.

Note: If the Macintosh or the EIII seem to repeatedly lock-up during the boot procedure, you may have duplicate SCSI ID's (see the SCSI ID Information Table). If nothing seems to function correctly, try booting all devices first, and then hooking up your EIII to Mac SCSI cable. Use this only as a last resort.

EIII Instrument Definition

The Edit Instrument dialog box is where you provide the necessary EIII communication information. Although the EIII can be set to any SCSI ID#, its default is ID number 6. This may be a good starting point, although you may change it at will. The SCSI ID parameter in the Edit instrument dialog box and the EIII's SCSI ID must always match.

To check the EIII itself for its current SCSI ID# press the Master button the EIII front panel followed by Disk Utilities.

Sync Mode

When running Alchemy with the EIII, a special mode option appears in the Network menu called "Sync Mode." In sync mode, when you execute an Alchemy editing or processing command, the affected part of the sound will automatically be transferred to the EIII. The result of this is that the sound in the Macintosh and the sound in the EIII stay in sync (are the same), without having to choose the Send Sound command. This type of operation is not practical without the high-speed benefit of SCSI data transfers.

The function of the Speaker Icon in the Alchemy function palette also changes when sync mode is switched on. In sync mode, clicking on the speaker icon (or depressing the space bar on the Macintosh keyboard) previews the sound you're editing directly out the EIII in 16 bit stereo. The Macintosh audio port will operate as explained in the Alchemy manual.

Sync mode can only remain active when the sound you're working on in Alchemy still resides in the memory of the EIII. For example, if you load a number of sounds into the EIII and transfer one over to Alchemy for editing, a link between the Macintosh and EIII is established, and sync mode will be operative. If you then erase that bank or load a new sound into the EIII, the link to the sound in Alchemy will be severed, and sync mode will be inoperative.

If you wish to isolate the editing performed on the Macintosh from the EIII, you can turn sync mode off by choosing it again on the Network menu.

A Word About SCSI ID's

For the EIII and Alchemy to boot-up and operate properly, you must be certain that every device in your SCSI chain has a unique SCSI ID#. Available SCSI addresses range from 0 to 7. Most SCSI hard disks and peripheral manufacturers allow you to select SCSI ID numbers manually to avoid potential conflicts. The table on the following page shows a list of SCSI ID numbers which may exist in your setup. Difficult or unmodifiable devices are followed by an asterisk.

SCSI Device ID Information Table

DEVICE	SCSI ID#
Macintosh internal hard disks	0*
Macintosh CPU's	7*
Macintosh external hard disks	selectable
Apple CSC CD-ROM	selectable
Emulator III internal hard disk	1*
Emulator III	6 (selectable)
Emu HD 300	4 (selectable)
Dyaxis processor	1*
Dyaxis hard disk	2*

General SCSI Notes:

- Make sure that no two SCSI devices have the same ID number.
- Make sure that your SCSI network is terminated correctly.
- The *sum total* length of all cables in a SCSI network must not exceed 15 feet.

*Difficult or unmodifiable devices.

IMS Dyaxis Specifics

Compatible software Version:
(All MacMix™ required).

Communications Types:
SCSI.

General Information:

- The first time a Dyaxis file is opened using Alchemy, it may take up to a minute for the waveforms to appear in a waveform window. This delay is a result of the fact that an amplitude envelope for the whole sound must be built before it can be viewed. Once a sound file has been viewed and saved a single time, the amplitude envelope will be retained, and the delay will no longer occur.
- When you edit a Dyaxis file, you are performing a disk-based operation, so you are actually changing the original file. For this reason it is always a good idea to edit only copies of your original Dyaxis files.
- Because Dyaxis files usually contain so much waveform data, you cannot use Cut or Inert commands when editing them. Commands which require that the entire waveform be shifted to the left or right are very time consuming with larger files, and other methods are more efficient.
- The RAM cache which you can set with the Edit Options... command on the Edit menu is Dyaxis-specific. The size of the RAM cache defines the largest waveform range which may be selected in a Dyaxis file. If you attempt to select a range that requires more memory than the RAM cache, the selection range will be adjusted to the RAM cache size.
- All recording and playing of sampled sounds on the Dyaxis system must be accomplished using the MacMix™ program that comes with the machine.
- Dyaxis files do not store Alchemy view memories.

A Note on MIDI Patch Bays

If you use a MIDI patch bay in your sampling network, it is important to use Alchemy's New Instrument dialog box correctly. Here are some facts you should know:

- The Patcher Channel is the MIDI channel which addresses and controls the patch bay. For some patchers (such as 360 Systems MIDI Patcher), the patcher channel is fixed (channel 16). For others (such as J. L. Cooper's MSB+), the patcher channel is adjustable. Consult your individual patcher's manual for more information.
- The Patcher Program is the actual patcher memory number which allows access to the sampler you are adding to your network. The patcher program is always described in a decimal system, so with patch bays such as the MSB+, which use a different numbering scheme you will need to adjust the program number to its decimal equivalent. All non-decimal patch bays contain a conversion table in their documentation.
- Many MIDI patch bays have a required input port. In other words, you need to plug the MIDI output from your Macintosh into a particular MIDI input on the patcher. In order for your MIDI patch bay to function, you will need to make sure that the correct input port is being used. Consult your patch bay's manual for more information.

PART NO. 01016-81030

Alchemy Manual Rev. C
01016-81030

©1996 Passport Designs, Inc.

All Rights Reserved

The following text describes the features and changes for Alchemy v3.0. The topics are:

- *Expanded Tool Palette
- *Multi-Tap Digital Delay
- *.WAV File Support
- *Apple Sound Manager 3.0 Support
- *Colors
- *PowerBooks and Alchemy: MIDI Manager version 2.0.2
- *OMS Support
- *ASR-10 Support
- *Force Instrument
- *Time Scale - Grain Size
- *Recording Sound
- *Rename Preset
- *DigiDesign Pro Tools Support
- *Incompatibilities
- *Manual Corrections

Expanded Tool Palette

Blending Controls

There is now a button in the Tool palette to turn Alchemy's Blending feature (Edit menu) on and off. Turning it on or off in the palette will turn it on or off in the menu and vice versa. There is also a numerical display to show the current blend time (in seconds). Clicking on the blend time display opens the Edit Options dialog with the Edit Blend Time field highlighted. See page 206 of the Alchemy Owner's Manual for more information about Blending.

Tool Help

As you move the arrow pointer over the Tool palette, the display at the bottom of the palette shows the name of the tool you're pointing to.

Multi-Tap Digital Delay

The Echo item in the Process menu brings up the Multi-Tap Digital Delay dialog. This allows you to add simple or more complex echo effects to the selected data.

The upper portion of the dialog displays one of 5 identical panels, one for each of the 5 taps. Use the Tap pop-up menu to select a panel. The enabled taps have a black diamond next to them in the pop-up. To enable additional taps (up to 5), choose an unused tap from the pop-up and, when its panel appears, check the Enabled checkbox.

Each tap can have its own delay time, initial delay, decay time, level, and pan position. Pan position is only available when you're processing a stereo soundfile and both channels are selected.

Tip: To select both channels, hold the mouse button down and drag the I-beam pointer along the line separating the two channels in the Waveform window. You can also double-click on the line to Select All.

The initial delay is used to set the time between the first occurrence of a sound and the first reflection (or first repeat) of the sound. The delay time is the time between subsequent repeats. The decay time determines how long the delayed signal will repeat after the first reflection occurs.

The relative amplitudes of the dry and wet signals can be adjusted in two ways. The Level control

sets the level of the processed signal for each tap. The Mix control sets the overall mix of dry and wet signals.

Tip: Hold [option] and click the thumb of the Level, Pan, or Mix control to center it. Centering the Mix control sets it to a 1:1 mixture of dry and wet signal.

Alchemy is shipped with a few preset delay settings, but you can also save an unlimited number of your own. To choose a preset, click and hold on the pop-up menu in the lower part of the dialog and choose the desired preset. To create a new preset, make your delay settings and click the New button. A dialog appears that enables you to name your preset and add it to the pop-up. The Store button allows you to save changes to an edited preset. The Delete button removes the current preset from the pop-up menu.

Given the nature of an echo effect, it is logical that the effect would extend beyond the selected range or, if the entire soundfile is selected, lengthen the sample. You can, however, force the effect to be constrained to the selected region. Simply check the Echo within selection only checkbox.

Parameters:

Delay time: 1 to 3000 milliseconds

Initial delay: 1 to 3000 milliseconds

Decay time: 0.0 to 30.0 seconds

Level: 0 to 127

Pan position: 0 (hard left) to 127 (hard right)

Mix: 0 (dry) to 127 (wet)

.WAV File Support

Alchemy 3.0 has the ability to read and write 8- and 16-bit .WAV files. Files translated with Apple File Exchange, PC Exchange, or similar programs can be opened in Alchemy if they have the type "WAVE" or a .WAV file extension. 8-bit files are converted to 16-bit when they are opened. The Save As dialog allows you to save files in either 8- or 16-bit .WAV format.

Apple Sound Manager 3.0 Support

Alchemy 3.0 supports version 3.0 of the Apple Sound Manager. That means that PowerMacs, AV Quadras, and add-on hardware that supports Sound Manager 3.0 can play back 16-bit audio directly. The Listen button in the Open Special dialog will also play loops in looped samples.

Colors

The Colors item in the Action menu lets you customize the color scheme in Alchemy's data windows. Choosing Colors opens the Color Selection dialog. This dialog contains a picture representing the Waveform window. Clicking on any of the various parts of the picture brings up a color picker that allows you to set the color for that part of the window. You can choose colors for the waveform, the waveform background, the ruler, the overview waveform, and the overview background. The colors in the Color Selection dialog will change to reflect your choices.

Note: The colors are saved with preferences (File menu).

PowerBooks and Alchemy: MIDI Manager version 2.0.2

Alchemy's installation program will automatically install MIDI Manager v2.0.2 if you are running on a PowerBook. MIDI Manager v2.0.2 will prevent the loss of incoming data that can occur when using one of these computers and transmitting large blocks of data such as System Exclusive transfers.

The fixes in MIDI Manager v2.0.2 only affect the modem port on PowerBooks. The printer port's behavior has not been improved and should not be used. MIDI Manager v2.0.1 is installed for all

other computers.

Note:

If at any time you wish to install MIDI Manager v2.0.2 you can do so using Alchemy's Custom installation option. You should be aware, however, that MIDI Manager v2.0.2 may cause problems during floppy disk insertion on the Macintosh Plus. According to Apple Computer, the fixes for the modem port are only required for the 140, 145, 160, 170, or 180 model PowerBooks.

OMS Support

If you are using OMS in your MIDI setup, you can choose OMS as your MIDI driver in Alchemy's Preferences dialog (Action menu). This affects the Alchemy interface in two instances. The Thru Settings and Instrument (New and Edit) dialogs will display an OMS Device pop-up menu that lists all the devices in your current OMS setup.

For more information, refer to the documentation that accompanies OMS.

ASR-10 Support

A new instrument type for the Ensoniq ASR-10 has been added to Alchemy's instrument list. The ASR-10 choice uses a specific protocol that includes sending of keymap and layer information to and from the ASR-10 via either MIDI or SCSI.

SCSI implementation in the ASR-10 is different from the SCSI implementation of other samplers supported by Alchemy. The ASR-10 provides power to SCSI bus and this can result in problems when connected to a Macintosh. By design, only one device is supposed to provide power to a SCSI chain. Since the Macintosh also provides power to SCSI, the ASR-10's attempt to control the SCSI chain can conflict with the Macintosh and result in a crash.

To circumvent this problem we have found it necessary to follow a specific "boot sequence" to connect an ASR-10 and a Macintosh together via SCSI.

Note:

Because the ASR-10's SCSI implementation is unique, Passport makes no claim that SCSI communication between Alchemy and an ASR-10 will work on all computers or in all situations. For assistance in trouble-shooting SCSI communication problems with an ASR-10 you should call Ensoniq for support.

Recommendations for SCSI Communication with an ASR-10

The ASR-10 uses SCSI ID#3 by default. This ID cannot be changed and before using your ASR-10 with your Macintosh you should first check to make sure no other device in your SCSI chain is using ID#3. In addition to external hard drives, be sure to also consider any internal devices such as CD-ROM drives. Many of the more recent Macintosh models come configured with an internal CD-ROM drive and these drives are often assigned to SCSI ID#3. You will have to either disconnect the CD-ROM drive within your Mac or change the SCSI ID for the CD-ROM unit before you can continue.

Direct Connection between an ASR-10 is possible if a 25 pin to 25 pin SCSI cable is used. In our experience, the "boot sequence" is extremely important for this connection to work. Connect your Macintosh to your ASR-10 with both units turned off. Turn on, "boot", the ASR-10 first. After the ASR-10 is up and running, boot your Macintosh. Finally, once the Macintosh is running, reboot the ASR-10. Variations on this procedure can be tried if you are unsuccessful, but these steps have yielded the most reliable results in our testing.

As 25 to 25 pin SCSI cables are sometimes difficult to obtain and a direct connection to a Macintosh can be tricky, the best solution is to place a SCSI device formatted for the ASR-10 between the Macintosh and the ASR-10. In this instance, all devices are connected with the power turned off, checking first for possible SCSI ID conflicts. The volume formatted for the ASR-10 (usually a hard drive or removable cartridge drive) is booted next. After the volume formatted for the ASR-10 has

powered up, boot the ASR-10. In this situation the ASR-10 will be loading its operating system from the hard drive or cartridge drive. After the ASR-10 has powered up, boot any additional, external SCSI drives before turning on the Macintosh.

Finally, before attempting to send or receive using SCSI with your ASR-10, you must also connect MIDI in and out from the ASR-10 to a MIDI interface connected to your Macintosh and configure the correct port and driver for MIDI communication in Alchemy. Although the ASR-10 will send and receive samples over SCSI, the keymap information and "handshakes" necessary to send it are transmitted via MIDI.

Force Instrument

Alchemy's Action menu contains a new item called "Force Instrument." This item displays a check mark when it's enabled. and can be enabled using the shortcut of the command and equals keys. Force Instrument applies the current selected instrument in the Network menu to each file as it is opened, overriding any instrument choice that may have been saved with the file in a previous Alchemy session.

Keyboard Equivalent: [command]+[=]

In addition, holding down the [option] key while choosing an instrument from the Network menu will set all currently open files to that instrument.

Time Scaling

A field has been added to the Time Scale dialog. In addition to the time scale factor, a "grain size" factor can be altered for more flexible results. Grain size defaults to a value of 30.

The grain size and default value have always been used for time scale operations in Alchemy but, prior to version 3.0, the grain size was inaccessible to the end user. It is now possible to decrease the grain size to as small a value as "2" or as great a value as the sample rate, size, and time scale factor will allow.

Grain size is an arbitrary unit used to define a "time slice." During a time scale operation, time slices are analyzed for frequency and amplitude content and a new time duration is constructed one slice at a time. Smaller grain sizes can result in a "smoother" end result, but there are trade-offs. To start with, smaller grain sizes will require more time for the computation. In addition, when using smaller grain sizes, the low frequency content may not be properly analyzed. The lower the frequency the more samples required to determine the frequency and so, if you use too small a grain size, the low frequency content in your sample may be lost entirely.

Using a larger grain size has the result of producing sounds that might best be described as more "chunky." The end result can sound similar to an echo effect as longer "time slices" of the sample are reproduced. Another nice bonus is that time scale operations are faster with larger grain sizes.

As with other DSP effects, the best way to understand and learn what does and doesn't work is to experiment. Rather than limit the use of this feature (by limiting the range), the field will accept anything from 2 on up and only warn you if the grain size or time scale amount is unusable. The grain size cannot be "larger" than the duration selected, for instance.

If you really must know...

As far as what a "grain" actually is, well...do you really want to know? Oh, all right. Basically, a grain is $1/1000$ of the sample rate plus a "taper" factor. The taper amount ensures the size result is an even multiple and is also used to blend the time slices together. Again, don't try to analyze this one too much - just fool around and see what works.

Recording Sound

Sound files can now be recorded directly into Alchemy using the Apple Sound Manager v3.0. The "Record Sound" item is in Alchemy's Action menu. Sounds are recorded into free memory (the amount of free memory available to Alchemy is displayed in the Tool Palette). A Sound Input device must be present and selected in the Sound In panel of your Mac's Sound Control Panel.

Setup

To begin, choose "Record Sound" from the Action menu. The "Record Sound" dialog will appear.

If you have more than one input device installed in your system, you can choose the device you wish to use from the device pop-up menu. The Sample Rate pop-up menu will reflect the sample rates available for the input device. Sample size and channel selection will also depend on the input device.

When the "Monitor" checkbox is enabled, the incoming sound signal will be sent through the Sound Manager to the selected output device. (Note: Monitor sends the sound to the Sound Output device selected in the Sound Manager, not the currently selected Audio Output selected in Alchemy). The slider to the right of the Monitor option controls the volume of the monitored signal.

Sound input levels are displayed using the horizontal meter bar near the center of the dialog. If you're running Alchemy on a color system, the input levels are color coded. Volume levels below maximum are displayed in blue. Levels that exceed the input device are shown in red. As input devices and content can vary, it is recommended that you perform some tests before you begin the actual recording. In general, the best recordings are obtained when signals are as loud as possible without distorting. After making a sample recording, you can check it in Alchemy's waveform window to see if the loudest portions are getting "clipped" or if the levels are correct (or too low). If you've ever heard a clipped digital recording, then you already know that it's a particularly obvious and ugly type of distortion and is in no way similar to "saturating" an analog recording tape.

Recording

Recording is begun by clicking on the Record dialog's Record button. Click Stop to finish recording. If you run out of RAM for recording, Alchemy will stop the recording process for you.

Only one file is created each time you open the Record dialog. Each time you click Record during the same session (that is, without exiting the dialog), a new recording is begun, replacing the previous "take." Click Play to hear the results of the previous record operation.

Click Done to exit the Record dialog. If you have made a recording, the newly recorded sound file will appear in an untitled window in Alchemy. To save the sound file use the standard Save As dialog and select a file format.

Record Sound and DigiDesign cards

To use a DigiDesign card for recording in Alchemy's "Record Sound" dialog, you must install the

appropriate DigiDesign Sound Driver extension for Sound Manager v3.0. This extension is NOT the same as the DigiSystem Init. Sound Drivers are used with the new Sound control panel device (CDEV) in conjunction with Sound Manager v3.0. To obtain a copy of the DigiDesign Sound Drivers contact DigiDesign.

Rename Preset

The Rename Preset option in the Send Sound dialog will only work with an E-mu Emax or a Korg T-series instrument.

DigiDesign Pro Tools Support

Alchemy has been updated to support the new Pro Tools III card. For Alchemy to recognize any of the DigiDesign cards or the MediaTime card there are two requirements. First, the DigiSystem Init must be loaded for all of the DigiDesign cards (the MediaTime card requires the RasterOps driver). Second, Alchemy will only recognize the first audio card it finds. First in this case means whichever card is in the lower numbered NuBus slot.

If you have a DigiDesign TDM card - also called "the Farm" - and Alchemy DOES NOT recognize and enable the Pro Tools option in the audio output menu, try swapping the TDM and Pro Tools card positions.

Incompatibilities

As of Alchemy's release, two older system extensions have been found to be incompatible with Alchemy v3.0. Both problems are related to new file procedures in version 3.0 that have been written specifically for system 7 and can result in a crash when using the Open menu.

Boomerang v3.0 from Now Software is an older version of the popular file utility now available as Super Boomerang. Version 4.0 of Super Boomerang is included in the Now Utilities software package. If you are still using the older version of Boomerang you will need to either exclude Alchemy from Boomerang's use or update to the latest version.

QuicKeys v2.0 from CE Software shipped with a variety of "extensions" found within the QuicKeys Preferences folder. One of these extensions - "Location" - has been found to be incompatible with Alchemy. Either remove the file from the extensions folder or contact CE Software for an update. (Note: the Location file is located in a folder called "extensions" which is in turn located within the QuicKeys folder within the Preferences folder! Don't confuse the QuicKeys "extensions" folder with the "extensions" folder in your system)

Manual Corrections

There is an error in the manual regarding a special Paste and Mix option in Alchemy. Under normal circumstances, when a range is selected on screen and you paste data from the clipboard, the data will be pasted from the beginning of the selection and will replace the existing data for the length of the clipboard data. If you hold down the option key and choose Paste, the clipboard data will be pasted from the end of the selection forward. In other words, you will align the end of the clipboard data with the end of the selected data. This also applies to the Mix command and can be used in either the Waveform window or the Harmonic window.

The error occurs on pages 5-30 and 5-34 of the old manual and pages 204 and 209 of the latest edition. The manual states that you should hold down the command key, not the option key.